

# **SANDIA REPORT**

SAND2017-10772

Unlimited Release

Printed September 2017

## **Developing strong, concurrent, multiphysics, multiscale coupling to understand the impact of microstructural mechanisms on the structural scale**

James W. Foulk III, Coleman N. Alleman, Alejandro Mota, Hojun Lim, David J. Littlewood, Guy Bergel, Steve Owen, Corey Ernst, Evdokia Popova, David Montes de Oca Zapiain, Surya Kalidindi

Prepared by

Sandia National Laboratories

Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Approved for public release; further dissemination unlimited.



**Sandia National Laboratories**

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology and Engineering Solutions of Sandia, LLC.

**NOTICE:** This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from  
U.S. Department of Energy  
Office of Scientific and Technical Information  
P.O. Box 62  
Oak Ridge, TN 37831

Telephone: (865) 576-8401  
Facsimile: (865) 576-5728  
E-Mail: [reports@adonis.osti.gov](mailto:reports@adonis.osti.gov)  
Online ordering: <http://www.osti.gov/bridge>

Available to the public from  
U.S. Department of Commerce  
National Technical Information Service  
5285 Port Royal Rd  
Springfield, VA 22161

Telephone: (800) 553-6847  
Facsimile: (703) 605-6900  
E-Mail: [orders@ntis.fedworld.gov](mailto:orders@ntis.fedworld.gov)  
Online ordering: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>





# **Developing strong, concurrent, multiphysics, multiscale coupling to understand the impact of microstructural mechanisms on the structural scale**

James W. Foulk III  
Coleman N. Alleman  
Alejandro Mota  
Guy Bergel  
Mechanics of Materials (8343)  
jwfoulk@sandia.gov

David J. Littlewood  
Multiscale Science (1444)  
Sandia National Laboratories  
djlittle@sandia.gov

Hojun Lim  
Computational Materials & Data Science (1814)  
Sandia National Laboratories  
hnlm@sandia.gov

Steven Owen  
Simulation Modeling (1543)  
Sandia National Laboratories  
sjowen@sandia.gov

Corey Ernst  
Elemental Technologies  
corey@elemtech.com

Evdokia Popova  
David Montes de Oca Zapiain  
Surya R. Kalidindi  
Woodruff School of Mechanical Engineering  
Georgia Institute of Technology  
surya.kalidindi@me.gatech.edu

## Abstract

The heterogeneity in mechanical fields introduced by microstructure plays a critical role in the localization of deformation. To resolve this incipient stage of failure, it is therefore necessary to incorporate microstructure with sufficient resolution. On the other hand, computational limitations make it infeasible to represent the microstructure in the entire domain at the component scale. In this study, the authors demonstrate the use of concurrent multi-scale modeling to incorporate explicit, finely resolved microstructure in a critical region while resolving the smoother mechanical fields outside this region with a coarser discretization to limit computational cost. The microstructural physics is modeled with a high-fidelity model that incorporates anisotropic crystal elasticity and rate-dependent crystal plasticity to simulate the behavior of a stainless steel alloy. The component-scale material behavior is treated with a lower fidelity model incorporating isotropic linear elasticity and rate-independent  $J_2$  plasticity. The microstructural and component scale subdomains are modeled concurrently, with coupling via the Schwarz alternating method, which solves boundary-value problems in each subdomain separately and transfers solution information between subdomains via Dirichlet boundary conditions.

Beyond cases studies in concurrent multiscale, we explore progress in crystal plasticity through modular designs, solution methodologies, model verification, and extensions to *Sierra/SM* and manycore applications. Advances in conformal microstructures having both hexahedral and tetrahedral workflows in Sculpt and Cubit are highlighted. A structure-property case study in two-phase metallic composites applies the Materials Knowledge System to local metrics for void evolution. Discussion includes lessons learned, future work, and a summary of funded efforts and proposed work. Finally, an appendix illustrates the need for two-way coupling through a single degree of freedom.

## **Acknowledgment**

We appreciate the patience and guidance of Jim Redmond, the project manager for this research. Jim leads multiscale efforts within ASC and it is our hope that the work completed during this LDRD will help provide the technical basis for multiscale modeling to impact our laboratory. The writers gratefully acknowledge the many fruitful discussions in the development of this work with Corbett Battaile, Irina Tezaur, Jakob Ostien, Fadi Abdeljawad, and Theron Rodgers.



## Contents

1	Introduction .....	11
1.1	Assertions .....	14
2	Concurrent multiscale: Methods .....	18
2.1	Microscale Physics .....	19
2.2	Concurrent multiscale modeling with the Schwarz Alternating Method .....	26
3	Concurrent multiscale: Findings .....	29
4	Progress in crystal plasticity .....	33
4.1	Viscoplasticity via multiple flow rules and hardening laws .....	34
4.2	Incorporating drag-based flow with nonlinear root-finding .....	35
4.3	Verification through quad-slip .....	40
4.4	Porting to the LAME material library .....	46
4.5	Extension to manycore architectures .....	48
4.6	Robustness suite .....	49
5	Advances in conformal microstructures .....	52
5.1	Application to a hexahedral workflow .....	52
5.2	Moving to a tetrahedral workflow .....	53
6	Structure-property linkages .....	60
6.1	MKS framework .....	60
6.2	Dataset generation .....	60
6.3	Reduced order model .....	64
6.4	Summary .....	66
7	Discussion and future work .....	68
7.1	Requisite components .....	68
7.2	A path forward through teaming .....	73
8	Conclusion .....	76
	References .....	78

## Appendix

A	Need for concurrent, two-way coupling .....	84
---	---	----

## Figures

1	A schematic illustrating concurrent coupling of the structural scale and the microstructural scale. Heterogeneous fields of plastic slip and hydrogen concentration at the microstructural scale transition to smooth fields at the structural scale. As these field evolve, the stresses induced by the high-pressure gas attempt to drive the crack forward. Substantial gradients exist at both scales. ....	13
2	Discovery of the dominant mechanisms that drive the performance of structures stems from efforts in microscale physics, strong multiphysics coupling, concurrent multiscale coupling, and a flexible research environment. ....	17

3	Microstructural domain, colored by grain ID. ....	21
4	Inverse pole figure with respect to the loading direction for a particular microstructural realization. ....	21
5	Load-displacement response of ten realizations of crystal orientations. ....	25
6	Reduced true stress-strain curves and the $J_2$ fit of ten realizations of crystal orientations. ....	26
7	Subdomains of tensile specimen. The domain is split into two grip subdomains, each colored blue, and a microstructural subdomain at the center of the gauge section, colored by grain number. The $50\mu\text{m}$ overlap regions can be seen at either end of the microstructural subdomain. ....	28
8	Contour of $F_{11}$ —the normal component of the deformation gradient in the loading direction. Localization is evident in the regions with red coloring. ....	29
9	(a) Reduction in cross-sectional area over time. Localization is evident in the vicinity of $-50\mu\text{m}$ , where the area is being preferentially reduced. (b) Rates of area reduction as a function of time. The signature of localization is seen in locations around $-50\mu\text{m}$ , as in (a). ....	30
10	Cross-sectional area profiles for simulations employing symmetry, fixed, and Schwarz boundary conditions. ....	31
11	Stress-strain response of microstructural domain: (a) Normal, along tensile axis; (b) Normal, perpendicular to tensile axis; (c) Shear in transverse plane. The solid lines with closed symbols correspond to the average response of the entire microstructural domain. The dashed lines correspond to the grain-averaged quantities. ....	32
12	A modular approach to crystal plasticity having the desired crystal physics and requisite numerics. ....	33
13	Comparing exponential hardening, saturation hardening, and Taylor hardening for tantalum single crystals of multiple orientations. Models were fit with experimental data in the $[100]$ orientation and then employed to predict the response in the $[110]$ and $[111]$ orientations. ....	35
14	Residual profile for a single slip system using (a) the power-law flow rule, and (b) the power-law with drag flow rule. ....	36
15	The incremental solution path for the power-law with drag residual of a single slip system using (a) Newton’s method, (b) trust region, and (c) back-tracking line search. ....	38
16	Comparison of (a) the elapsed model time for a given simulation step, and (b) the total material model iterations per time step using Newton’s method with and without drag, trust region, and back-tracking line search. ....	39
17	Profile of a single component of the power-law with drag residual using random combinations of slips for a 12 slip system material point. ....	39
18	The Rubik’s cube block setup. ....	40
19	The iterative solution and residual path for a representative material point in the Rubik’s cube problem using (a) power-law flow rule and Newton’s method, and (b) power-law with drag flow rule and Newton’s method. ....	41
20	Constitutive response (solid line) for a strain rate of $10^0 \text{ s}^{-1}$ . Model equations are (52) and (55), with parameters taken from Table 5. The dashed line is the saturated flow stress $\bar{S}_{11}$ calculated from Equation (60). ....	46

21	Speedup for the OpenMP parallel implementation of the $J_2$ and crystal plasticity models using Kokkos. (a) 1000 instances of the $J_2$ model. (b) 1,000,000 instances of the $J_2$ model. (c) 1000 instances of the crystal plasticity model. (d) 10,000 instances of the crystal plasticity model. ....	50
22	We developed a robustness suite to quantify and monitor the performance of our algorithms for problems of increasing complexity. ....	51
23	Polycrystalline microstructure generated from kinetic Monte Carlo grain growth simulations having 143 grains. (a) Full FEA mesh generated from $200 \times 100 \times 100$ grid, (b) close-up of mesh, (c) view of three grains in FEA mesh. ....	53
24	Polycrystalline microstructure generated from phase field grain growth simulations having 52 grains. (a) Phase field representation of the 3D grain microstructure. The scalar field $\phi$ is used for visualization, where (red) blue denote (GBs) bulk grains. (b) FE discretization of the microstructure and (c) view of three grains in FEA mesh	54
25	Electron backscatter diffraction (EBSD) measurement and FE discretization using interface conformal hexahedral mesh. ....	54
26	Tetrahedral workflow for conformal boundaries that leverages Sculpt, Cubit, and MeshGems. ....	55
27	Conformal grain boundaries reconstructed from a grain-growth simulation and discretized with higher-order tetrahedral elements subjected to substantial deformations. A far-field engineering strain of 0.45 was achieved in 225 time steps. The volume-averaged axial stress and grain-boundary “ridges” highlight the grain-grain interactions. ....	57
28	Illustrating the axial component of the total deformation gradient at and applied engineering strain of $\sim 0.6$ . The deformation localizes along grain boundaries with local stretches are on the order of 9. ....	58
29	Generalized MKS workflow for extracting structure-property linkage. ....	61
30	2-phase microstructure classes used in this study along with the relevant details of how they were generated. ....	63
31	CDF of damage parameter (top 2%) and contour plots of the top 2% of Damage parameter for Classes 6, 7, 8 and 9. Color scale bar is the same for all contour plots.	65
32	Top 2% of CDF of the original fit to GEVD (shaded area) and GEVD from predicted parameters (dashed lines) for all of the microstructures. ....	67
33	Schematic showing different levels of adaptive mesh refinement via recursive bisection. The resulting octree data structure provides increased accuracy to resolve grain boundary curvature and triple junctions. ....	71
34	Cases of substantial refinement that may stem from periodic boundary conditions in the grain growth simulations. More study is needed to understand the relevance of these features. ....	72
35	Components of a technical basis that enable research efforts, mature research environments, support production implementation, and ultimately serve our mission and foster collaboration with national/international communities. ....	75
A.1	Thought experiment to investigate the ramifications of one-way and two-way coupling. The figure is to scale with $L = 5L_f$ . ....	85
A.2	In one-way transfer, the coarse scale and the fine scale are not in equilibrium, $P_c \neq P_f$ . ....	85

- A.3 Two-way coupling localizes the displacement field in the fine scale. We remark that not only is one-way transfer inaccurate but the methodology is not conservative. 86

## Tables

1	Parameter values for grain growth simulation. ....	20
2	FCC slip systems. ....	23
3	Crystal plasticity parameters used in this study to model a representative stainless steel. ....	24
4	$J_2$ plasticity parameters fit to the reduced crystal plasticity simulations. ....	26
5	Crystal plasticity parameters employed for verification in uniaxial stress. ....	46



# 1 Introduction

In many applications of interest, the mechanical performance of metallic structures is driven by microstructure. Strongly coupled processes can localize grain deformations through heterogeneous plasticity and void nucleation. Phenomenological models of these processes are insufficient to observe, analyze, and understand problems such as the hydrogen embrittlement of austenitic stainless steels. The ability to make engineering design decisions depends critically on understanding and quantifying the effects of the processes, and thus on resolving microstructure within the “structural” scale of an engineered component. The dominant dissipative mechanisms for candidate face-centered cubic (FCC) metals result from intimately coupled processes evolving over many grains ( $\sim\mu\text{m}$ ) embedded within a structure ( $\sim\text{m}$ ) subjected to complex loading in non-inert environments. In some cases, strong multiphysics coupling and concurrent multiscale models are required to resolve failure processes. The simulation of failure may also require advanced material models coupling the phenomenology of crystal plasticity with void nucleation ( $\sim\text{nm}$ ) and growth.

Localization and the subsequent fracture/failure of engineering alloys intimately couples the structural scale driving forces derived from the mechanical action applied to a body and the resistance to fracture, which is governed by the metallurgical details of microstructure. Simply put, the global loads are trying to tear apart a local microstructure that struggles to accommodate, dissipate, and remain intact. The performance of the structure is quantified through global quantities, while the variability of the microstructure is quantified through local quantities. The ability to predictively connect local variability to global margins is the promise of multiscale modeling for engineering applications. However, the application of multiscale models to study failure remains a significant challenge. One cannot understand the impact of variability on localization through volume-averaged quantities alone [4, 19, 64]. Hierarchical multiscale methods such as first-order homogenization (volume averaging) are commonly employed in analyzing elastic and elasto-plastic behavior, but this type of analysis cannot be used to capture the energy dissipation of the structure that is due to local processes [12, 14]. There exist hierarchical methods to tie ensembles of microstructure to structural scale models to estimate the impact of microstructure [26, 39]. However, these methods are prohibitively expensive at the structural scale, and furthermore, it is not yet clear which features of polycrystal morphology are relevant, making it difficult to choose an appropriate ensemble.

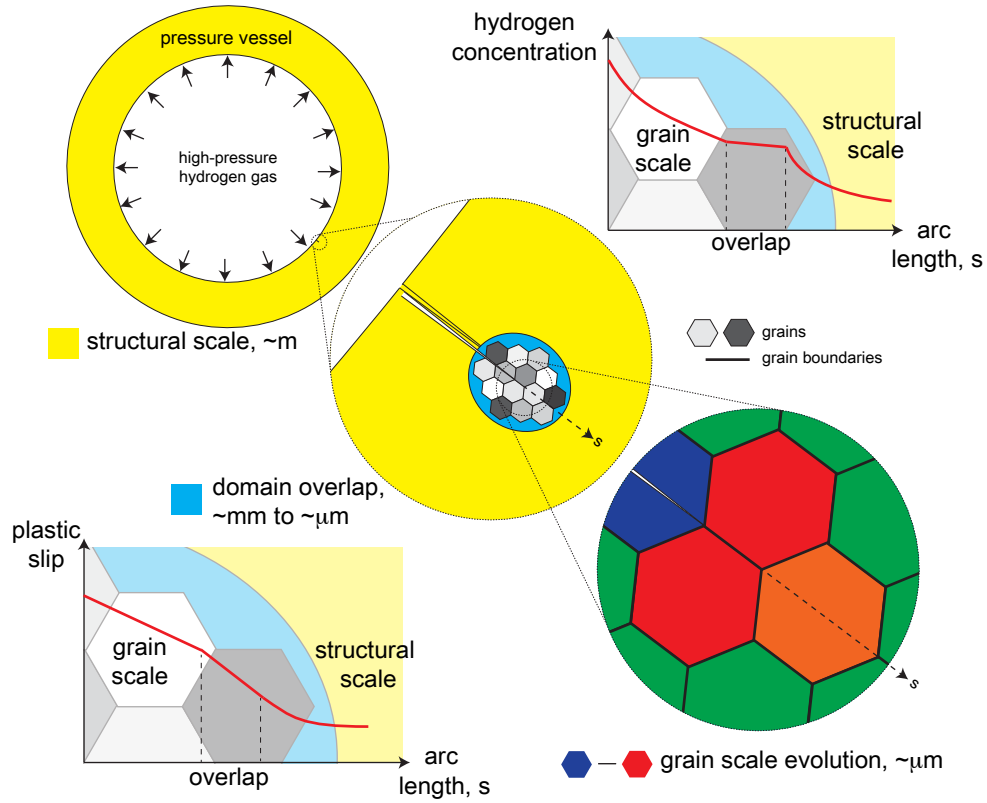
In determining the effective elastic properties of a polycrystalline aggregate, it is common practice to employ studies of representative volume elements (RVEs). This practice is predicated on the success of such analyses that were demonstrated originally for the effective properties of composites assumed to have a periodic microstructure of a regularly-shaped second phase embedded in an isotropic primary phase [31]. Application of this type of analysis to polycrystalline microstructures is effective for elastic properties even where the assumption of periodicity is not strictly satisfied [50]. Successful extension of this approach to modeling metal plasticity within the regime of uniform elongation has also been demonstrated, for example in [38], but for analyzing necking behavior it is not clear that this approach is tenable. In fact, in this study it is shown that the influence

of boundary conditions on the onset of necking is too strong to allow meaningful study of this phenomenon with RVE-type analyses. On the other hand, *concurrent* multiscale modeling provides the means to subject microstructure to the (boundary) conditions that would naturally arise in the loading of a structure. In studying localization, concurrent models can focus computational power in the vicinity of necking, while phenomenological constitutive models can be employed in regions of lesser interest. While these reduced order models are insufficient to analyze the microstructural influence on localization, it is hypothesized here that they are sufficient for transferring loads the structural scale in the form of more realistic boundary conditions on the region of interest. These models are significantly less expensive to compute than a physically-based microstructural model such as a continuum crystal plasticity model, so that the additional computational cost of modeling the entire structure can be afforded. The most straightforward incarnation of this type of model avoids invoking any phenomenology in the hierarchical passage of information across scales and instead couples the values of the primary solution variable between different parts of the simulation domain. In the case of the present study, the focus is on mechanics and the balance of linear momentum, so the displacement field is the coupled quantity. The coupling of the displacement on interfaced [13, 25] or overlapping [51] subdomains enables the flow of information between scales, while the internal state variables are maintained locally and do not require transfer. This type of concurrent coupling can be accomplished through direct methods such as is used here, where both scales are solved simultaneously, or methods which iterate between the microscale and the structural scale.

A schematic of concurrent coupling is illustrated in Figure 1. Consider a pressure vessel with a manufactured flaw that is exposed to high-pressure hydrogen gas. Microstructures are embedded at the crack tip and loaded through concurrent coupling. Hydrogen transport [22, 81] is operable at both scales and substantial gradients exist in both mechanics (plastic slip) and transport (hydrogen concentration). In fact, the localization of the deformation may occur via deformation bands, twins, and planar phase transformations [3, 52, 54, 75]. Although current approaches can capture the heterogeneity in the solution fields due to a geometric discontinuity, we cannot capture microstructural heterogeneity without concurrent multiscale coupling. Our ability to predict void nucleation as a function of microstructure, gas pressure, and time will characterize the performance of the structure. Although current microstructural models do not admit stress-softening and the loss of ellipticity [20], future efforts will capture void nucleation and regularized damage evolution [37].

Given the challenges involved in spatial multiscale modeling, one wonders if joining the structural scale and the microstructural scale via refinement along a contiguously meshed interface is a viable option. Although this is a natural leaning, there are major limitations:

1. There are limits to mesh refinement. The eigenvalues of the tangent stiffness scale with the element size. Elements in resolved microstructures ( $\sim 10$  nm) may be 6 orders of magnitude smaller than elements at the structural scale ( $\sim 1$  cm). This results in a highly ill-conditioned numerical system that is difficult to solve. Iterative multiscale methods are not subject to these constraints.



**Figure 1.** A schematic illustrating concurrent coupling of the structural scale and the microstructural scale. Heterogeneous fields of plastic slip and hydrogen concentration at the microstructural scale transition to smooth fields at the structural scale. As these field evolve, the stresses induced by the high-pressure gas attempt to drive the crack forward. Substantial gradients exist at both scales.

2. The transition at the interface of the structural and microstructural subdomains is abrupt. By contrast, classes of multiscale methods employ a domain overlap that may provide a smoother transition between constitutive models.
3. For explicit transient dynamics, element sizes at the microstructural scale control the global time step. If the coarse and fine scale are spatially separable, one can also partition the time integrator. The structural scale can be explicit while the microstructural scale can be implicit. Tractable and unconditionally stable time steps can be achieved at the microstructural scale.

Through ensembles of microstructures, we will be able to probe localization and subsequent failure processes to determine the performance of engineering structures. As a first step toward the study of localization and void nucleation, the current work will only focus on the balance of linear momentum and developing a framework for studying the effects of crystal anisotropy on localization. This analysis focuses exclusively on the role of polycrystalline microstructure in developing heterogeneous mechanical fields that will eventually result in strain localization.

## 1.1 Assertions

Before venturing into the body of the document, we believe it helpful to enumerate a few assertions that motivated the original proposal. Major components are illustrated in Figure 2. We are motivated to discover the dominant mechanisms that drive the performance of structures. We believe:

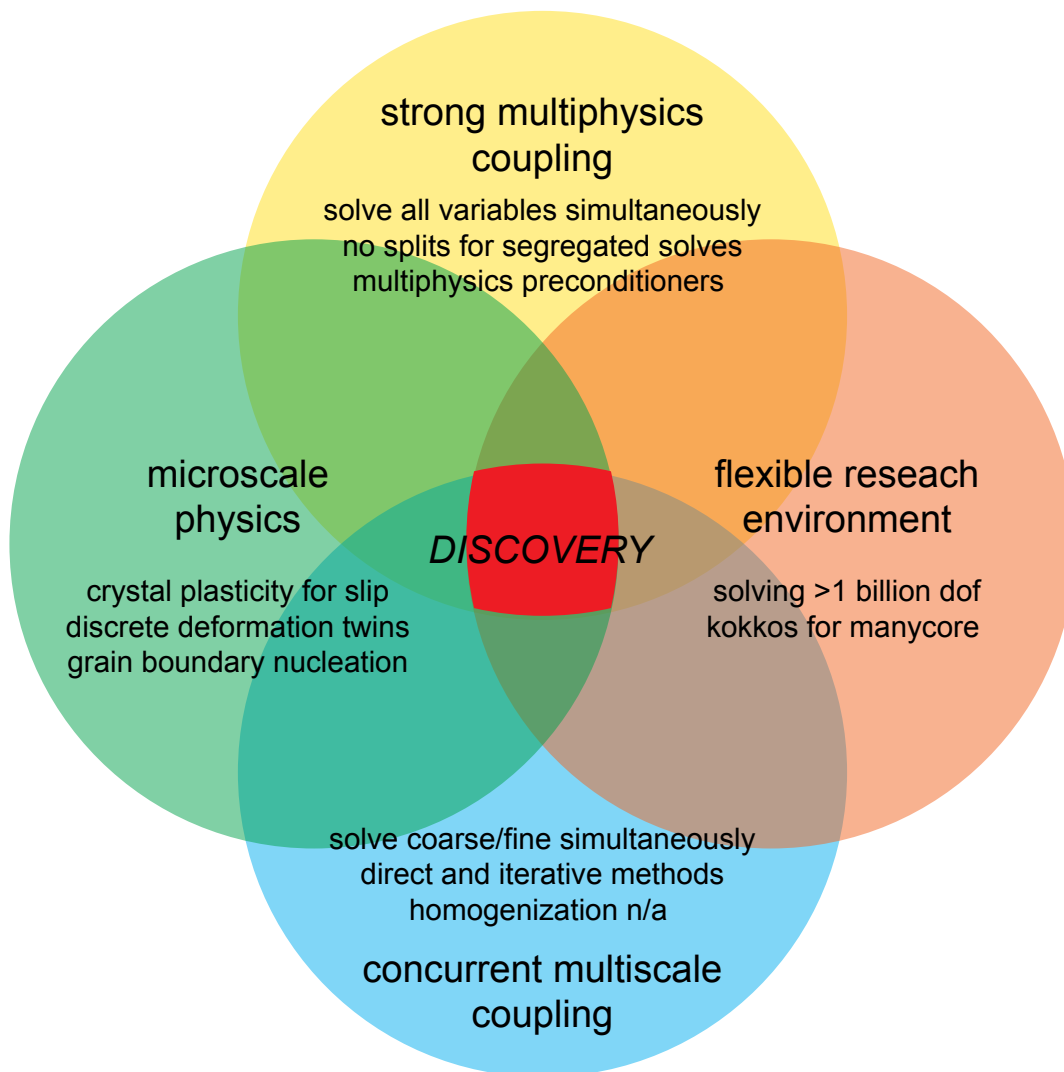
1. SNL applications invoke microstructure. This work was motivated through collaboration with component groups which have design specifications for microstructure.
  - (a) Resolve microscale physics. We seek convergent spatial and temporal descriptions. We believe that these processes can be achieved through refined meshes and implicit time integration.
  - (b) Honor heterogeneous microstructures. Although voxelated descriptions may have utility for far-field metrics, conformal grain boundaries are needed to resolve triple-junctions and grain-grain misorientations. If we believe that voids nucleate from grain boundaries, we cannot convolute the singularities that stem from voxelation with strain localization that may stem from the interaction of multiple grains along the load path.
2. Connecting scales provides a means for resolving behavior that may be difficult to deconvolute with phenomenological models and experimental findings. Multiscale coupling can also provide a stronger technical basis for quantifying performance metrics with increased confidence.
  - (a) Homogenization requisite. Homogenization is extremely helpful in understanding coarser level descriptions. In fact, we assert that although one is invoking additional physics at the microscale, those physics are fundamental and can be employed to yield complexities that may be difficult to describe and distinguish with phenomenology. For multiscale coupling in which different model descriptions are “connected” via an interface or

domain overlap, the impedance mismatch can be mitigated with homogenization of the fine scale. The limitation of first-order homogenization is the potential for bifurcation of the representative volume element.

- (b) Embedding a lower length scale with concurrent coupling. In cases dominated by localized deformations, we seek to resolve the fine scale and connect that scale to the coarse scale via concurrent coupling. Because the process of localization necessarily unloads the coarse scale, two-way coupling is needed. In this mindset, all localization must be resolved at the fine scale with appropriate methods for regularizing the boundary-value problem. There is not concept of a representative volume element. One just resolves localized regions with microstructure.
3. Flexible research environment. We assert that our development environment should both permit and encourage rapid prototyping of models and methods. This is, in fact, the reason why we invest in a development environment rather than implement methods in the production codes. We attempt to strike a balance between reusability and utility without tremendous complexity.
    - (a) Automatic differentiation. Automatic differentiation via Sacado provides infrastructure for the implicit solution of both the governing equations and the material model update. We assert that tools like automatic differentiation (AD) are helpful in making components modular to permit both the reduction and addition of microscale physics.
    - (b) Resolution through manycore/GPUs. This project initiated as Kokkos was being adopted for Advanced Technology Deployment and Mitigation (ATDM). We also sought to adopt Kokkos with the allure of a massive finite element simulation running on two GPU cards in your own workstation. Although we still retain this vision, the complexity of the venture has been sobering to our own efforts in the development environment.
  4. Honor reality through multiphysics coupling. Experiments are inherently multiphysics. Abnormal environments may require thermomechanical coupling. Hydrogen embrittlement requires chemomechanical coupling. We accept and assert that our most difficult problems involving localization are frequently multiphysics problems. We should be prepared to solve these classes of problems in both our development and production environments. In fact, the only way to truly know if the problem is not multiphysics is to indeed solve both the multiphysics and single physics problems and subsequently prove that the solutions yield the same quantities of interest.
  5. Enable property-structure linkages. When this project initiated, we sought to quantify the performance metrics at the structural scale through ensembles of microstructure. After the first year, Hojun Lim initiated a collaboration with Professor Kalidindi at Georgia Tech with the hope of applying Material Knowledge Systems (MKS) to microstructure. The complexities drove idealized representations but the notion of relating statistical descriptions of geometry to statics descriptions of performance remained. We assert the value of this work with a future focus on path-dependent processes.
  6. Better to stay in the trenches until it works. We are not asserting that our team takes the moral high ground with regard to implementation. We only seek to stress that developing truly

robust methods reveals fundamental knowledge of the problem. We seek that fundamental knowledge and are quite willing to reduce scope rather than retain scope to obtain a “story” without understanding. This path necessarily involves a greater focus on verification, testing, and achieving robustness.

Because this work stems from a Laboratory Directed Research and Development (LDRD) project, the document is organized into sections that reflect our publications, interests, and collaborations. Section 2 and Section 3 review both the methods and findings that stem from concurrently coupling the microstructural scale to the structural scale [5]. In Section 4, we review the modular nature of our approach with a specific focus on solution methods, new architectures, and the formulation of a robustness test suite. Section 5 seeks to provide the reader with a review of the technologies we are applying to generate conformal microstructures. Findings for conformal grain boundaries are also presented. Section 6 presents key findings from the collaboration with Professor Kalidindi at Georgia Tech to develop property-structure linkages [69]. Section 7 attempts to place the work in perspective with lessons learned and an outline for future work, funding, and collaboration. Finally, we summarize our contributions in Section 8.



**Figure 2.** Discovery of the dominant mechanisms that drive the performance of structures stems from efforts in microscale physics, strong multiphysics coupling, concurrent multiscale coupling, and a flexible research environment.

## 2 Concurrent multiscale: Methods

This study seeks to develop and exercise a subset of the computational tools requisite for discovery of conditions conducive to localization. The ultimate goal of this work is to use physically based models to discover features in the phenomenology of localization during necking and correlate those features with morphological characteristics of polycrystal microstructure. It is asserted that an understanding of this process cannot be gained through the use of phenomenological models that presuppose knowledge about the behavior to be studied. In an effort to generalize needs beyond the current study, we propose several broad requirements of a simulation framework and models that must be satisfied in order to pursue the discovery process. First, we assert that the applicability of kinematics and constitutive behavior based on the assumption of infinitesimal deformations can only be understood in the context of the more general finite deformation theory. Even where the end goal is a homogenized model that employs assumptions of infinitesimal strain, it is understood that local deformations are heterogeneous and may not admit this assumption. To that end, in this study, a rigorous treatment of the microscale physics under finite deformations is implemented in order to gain insight into the fundamental mechanics of localization and to determine the appropriate set of assumptions for reduced-order models.

Second, for the class of problems illustrated in Figure 1, multiphysics coupling may indeed be requisite to understand the interplay between mechanisms of deformation and diffusion operating at different spatial and temporal scales. Although we do not address multiphysics coupling in the current work, we solve all the field variables simultaneously [78, 85] through “strong” coupling. The methodology does not rely on problematic operator splits [8, 9] and the corresponding segregated schemes. Third, if we accept that the importance of microscale physics and multiscale coupling on the failure process, we must facilitate the two-way passage of information from the microscale to the structural scale. In this work, we invoke the Schwarz alternating method [51, 76] to tie the microscale to the structural scale through the primary fields. This avoids the invocation of scale separation and first-order homogenization in the critical region of interest. Specifically, this framework does not rely on the identification of metrics that can be “passed” from one scale to another. Instead, the Schwarz alternating method is invoked for solving the same governing equations at two scales concurrently through alternating Dirichlet boundary conditions. We remark that the extension of the Schwarz alternating method to multiphysics is straightforward and will be the subject of future research.

Lastly, on the computational side, the solution of these large, complex problems requires a scalable, flexible environment to enable deployment of new physics and solution schemes. This project leverages an open-source code base, Albany [74], for the development environment. Within this general finite element partial differential equation (PDE) solution environment, the Laboratory for Computational Mechanics (LCM) [80] houses the governing equations and constitutive models for finite deformation solid mechanics. This environment takes advantage of the Trilinos library [30, 66, 67] and the accompanying packages for elements (Intrepid), multiphysics (Phalanx), automatic differentiation (Sacado), and solvers (NOX). In addition, the Kokkos package [21] is



leveraged to provide manycore performance portability and execute the multiphysics, multiscale framework on multicore central processing units (CPUs) and graphical processing units (GPUs). Initial studies on the scalability of the crystal plasticity model for manycore performance under Kokkos confirm that the current software functions as expected in a multithreaded mode, but those findings are outside the scope of what is presented here.

This manuscript is partitioned into sections devoted to microscale physics and concurrent multiscale coupling with application to the study of localized plastic deformation of a stainless steel tensile specimen. First, the microscale physics are fully resolved; microstructure is created by a physically driven grain growth model and the mechanical properties are analyzed with an anisotropic crystal elasto-viscoplasticity model. Second, these microscale results are upscaled using homogenization to obtain parameter values for the far-field macroscale constitutive relations. Last, the microscale and macroscale regions are coupled together and simulated in a concurrent multiscale framework to study the transfer of load from the grips of a tensile specimen into the gauge section, where deformation localizes during necking.

## 2.1 Microscale Physics

To fully resolve the potential mechanisms at play in the problem illustrated in Figure 1, it is necessary to model ensembles of microstructures with discrete representations of grains. In that vein, this work employs crystal plasticity (CP) [23, 42, 45] to model intragranular deformations for FCC grains, while postponing the inclusion of the effects of grain boundaries to future work.

### Computational microstructure

The polycrystal microstructure used in this study was created using the classical Monte Carlo Potts grain growth model [33, 72] in the Stochastic Parallel PARTicle Kinetic Simulator (SPPARKS) developed at Sandia National Laboratories. A more detailed description of the model can be found in [33, 72], but a brief explanation is given here.

The Monte Carlo Potts model evolves spins (here, grain IDs) on a fixed spatial grid based on a grain boundary energy  $E$  and a probability  $P$  of accepting a change in a neighboring grain ID. This evolution represents the growth of existing grains through consumption of neighboring grains by selectively reassigning the grain IDs of grid points adjacent to a grain boundary. The probability of changing the grain ID at each grid point in a given time step is

$$P = \begin{cases} \exp\left(\frac{-\Delta E}{k_B T}\right) & \Delta E > 0 \\ 1 & \Delta E \leq 0 \end{cases} \quad (1)$$

where  $\Delta E$  is the change in energy associated with the potential grain ID change,  $k_B$  is Boltzmann's

constant, and  $T$  is the simulation temperature. The values of the parameters used in this study are given in Table 1. By biasing the probability of acceptance to promote lower energy states, the

**Table 1.** Parameter values for grain growth simulation.

$N$	$L$	$k_B T$
128000	26	0.5

evolution is toward lower total energy from lowering the grain boundary content through grain growth. The probability of acceptance is also governed in part by the simulation temperature, which is used here as a parameter to control the curvature of the grain boundaries [86]. It is assumed here that all grain boundaries introduce the same energy per unit surface area, so that the total grain boundary energy in a microstructure is represented by the sum of the contributions between neighboring lattice sites in different grains as

$$E = \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^L [1 - \delta(q_i, q_j)] \quad (2)$$


where  $N$  is the number of lattice sites in the simulation,  $L$  is the number of neighbors for each lattice site, and  $q_i$  is the grain ID of site  $i$ . From Equation (2), it can be seen that changing a grain ID of a point to match the grain IDs of more of its neighbors lowers the energy contribution of that point. As such, the trend is toward more neighboring grid points sharing the same grain ID (grains growing).

To create the microstructure for this study, 2000 grain IDs are initially assigned to random grid points within a 40 x 40 x 80 simulation cell, and the SPPARKS simulation is allowed to run until a microstructure with 151 grains is produced. The physical scale of the microstructure is set to  $100\mu\text{m} \times 100\mu\text{m} \times 200\mu\text{m}$ . The average grain size (equivalent sphere diameter) is thus approximately  $12.7\mu\text{m}$ . The morphology shown in Figure 3 is used for all of the simulations used in this study. An ensemble of microstructures is created by assigning a different uniform texture to each realization within the ensemble. The inverse pole figure for the grain orientations in one such realization are plotted in Figure 4 as an example.


### Single crystal constitutive equations

The kinematics of the crystal elasto-plasticity model is based on a well-established continuum formulation [23, 32, 40, 41, 68, 71] invoking a multiplicative decomposition of the total deformation gradient  $\mathbf{F}$  into elastic and plastic parts:

$$\mathbf{F} = \mathbf{F}_e \cdot \mathbf{F}_p \quad (3)$$



**Figure 3.** Microstructural domain, colored by grain ID.



**Figure 4.** Inverse pole figure with respect to the loading direction for a particular microstructural realization.

Here,  $\mathbf{F}_e$  and  $\mathbf{F}_p$  are the elastic and plastic part of the total deformation gradient, respectively. The elastic Lagrangian strain  $\mathbf{E}_e$  is given in terms of  $\mathbf{F}_e$  as

$$\mathbf{E}_e = \frac{1}{2} (\mathbf{F}_e^T \cdot \mathbf{F}_e - \mathbf{I}) = \frac{1}{2} (\mathbf{C}_e - \mathbf{I}) \quad (4)$$

where  $\mathbf{C}_e$  is the elastic right Cauchy-Green deformation tensor. The elastic constitutive law relates the second Piola-Kirchhoff stress in the intermediate configuration to the elastic Lagrangian strain as

$$\mathbf{S} = \mathbb{L} : \mathbf{E}_e \quad (5)$$

where  $\mathbb{L}$  is the elasticity tensor. Solving Equation (3) for  $\mathbf{F}_e$ ,

$$\mathbf{F}_e = \mathbf{F} \cdot \mathbf{F}_p^{-1} \quad (6)$$

it is seen that the stress  $\mathbf{S}$  can be obtained for a given  $\mathbf{F}$ , given a solution for  $\mathbf{F}_p$ . The plastic constitutive relations developed in the following section produce a set of equations to obtain  $\mathbf{F}_p$  for a given  $\mathbf{F}_e$ , forming a coupled set of equations. The solution to these equations is a split of  $\mathbf{F}$  into  $\mathbf{F}_e$  and  $\mathbf{F}_p$  such that Equation (3) holds.

The velocity gradient  $\mathbf{L}$  in the current configuration is written as:

$$\mathbf{L} = \mathbf{L}_e + \mathbf{L}_p = \dot{\mathbf{F}} \cdot \mathbf{F}^{-1} \quad (7)$$

where,  $\mathbf{L}_e$  and  $\mathbf{L}_p$  are elastic and plastic parts of the velocity gradient, respectively, and can be represented as follows:

$$\mathbf{L}_e = \dot{\mathbf{F}}_e \cdot \mathbf{F}_e^{-1} \quad \text{and} \quad \mathbf{L}_p = \mathbf{F}_e \cdot \mathbf{L}_p^* \cdot \mathbf{F}_e^{-1}. \quad (8)$$

Here, the plastic part of the velocity gradient in the intermediate configuration is given by

$$\mathbf{L}_p^* = \dot{\mathbf{F}}_p \cdot \mathbf{F}_p^{-1}. \quad (9)$$

As the crystal deforms, the lattice is stretched and rotated according to  $\mathbf{F}_e$ . Under the conventional assumption that the plastic part of the deformation gradient does not cause lattice rotation, the Schmid tensor  $\mathbf{P}_0^\alpha$  in the reference and intermediate configurations coincide, and are given by

$$\mathbf{P}_0^\alpha = \mathbf{s}_0^\alpha \otimes \mathbf{n}_0^\alpha. \quad (10)$$

**Table 2.** FCC slip systems.

$\alpha$	Slip System	$\alpha$	Slip System	$\alpha$	Slip System	$\alpha$	Slip System
1	(011)[111]	4	(011)[ $\bar{1}\bar{1}\bar{1}$ ]	7	(011)[ $\bar{1}\bar{1}\bar{1}$ ]	10	(011)[111]
2	( $\bar{1}$ 01)[111]	5	(101)[ $\bar{1}\bar{1}\bar{1}$ ]	8	( $\bar{1}$ 01)[ $\bar{1}\bar{1}\bar{1}$ ]	11	(101)[111]
3	(110)[111]	6	( $\bar{1}$ 10)[ $\bar{1}\bar{1}\bar{1}$ ]	9	(110)[ $\bar{1}\bar{1}\bar{1}$ ]	12	( $\bar{1}$ 10)[111]

Here,  $\mathbf{s}_0^\alpha$  and  $\mathbf{n}_0^\alpha$  are the initial unit vectors in the slip direction and the slip plane normal direction on the  $\alpha$ -th slip system, respectively. The twelve  $\{111\} \langle 110 \rangle$  FCC slip systems used in the present model are listed in Table 2. Assuming plastic deformation is caused by the dislocation slip, the plastic part of the velocity gradient can be written in the intermediate configuration as [68]:

$$\mathbf{L}_p^* = \sum_{\alpha} \dot{\gamma}^\alpha \mathbf{s}_0^\alpha \otimes \mathbf{n}_0^\alpha = \sum_{\alpha} \dot{\gamma}^\alpha \mathbf{P}_0^\alpha. \quad (11)$$

It is implied that the summation in  $\alpha$  accounts for all the slip systems. Combining Equations (9) and (11), a governing equation for  $\mathbf{F}_p$  is derived:

$$\dot{\mathbf{F}}_p = \mathbf{F}_p \cdot \mathbf{L}_p^* = \mathbf{F}_p \cdot \left( \sum_{\alpha} \dot{\gamma}^\alpha \mathbf{P}_0^\alpha \right) \quad (12)$$

It is noted that in this form, an  $\mathbf{F}_p$  that is initially isochoric ( $\det(\mathbf{F}_p) = 1$ ) will remain isochoric throughout the deformation. To preserve this feature of  $\mathbf{F}_p$ , the numerical integrator used in this study employs the exponential map for the update.

The critical aspect of the single crystal constitutive equations in the crystal plasticity model is the relation between the slip rate and the applied stress. One of the most widely used form for a viscoplasticity model is the power-law function [36]:

$$\dot{\gamma}^\alpha = \dot{\gamma}_0^\alpha \left| \frac{\tau^\alpha}{g^\alpha} \right|^{1/m} \text{sgn}(\tau^\alpha) \quad (13)$$

Here,  $\dot{\gamma}_0^\alpha$  is the reference shear rate and  $m$  is the rate sensitivity factor.  $\text{sgn}(\tau^\alpha)$  is the sign of the shear stress on slip system  $\alpha$ , which is always  $\pm 1$ . Assuming isochoric plasticity, the driving force  $\tau^\alpha$  for slip on system  $\alpha$  is

$$\tau^\alpha = \det(\mathbf{F}_e) \boldsymbol{\sigma} : \mathbf{P}^\alpha = \mathbf{C}_e : \mathbf{S} : \mathbf{P}_0^\alpha. \quad (14)$$

In the case of precipitation hardening, the hardness values  $g^\alpha$  evolve according to the following rule

$$\dot{g}^\alpha = \dot{g}_0 \left( \frac{g_s - g^\alpha}{g_s - g_0} \right) \sum_{\beta} 2 \left| \mathbf{S}_{ij}^\alpha \mathbf{S}_{ij}^\beta \right| |\dot{\gamma}^\beta|, \quad (15)$$

where  $\dot{g}_0$  is a hardening rate parameter,  $g_0$  is the initial hardness,  $g_s$  is the saturation hardness, and  $\mathbf{S}$  is the symmetric part of the Schmid tensor. The saturation hardness is given by

$$g_s = g_{s_0} \left| \frac{\dot{\gamma}}{\dot{\gamma}_0} \right|^\omega, \quad (16)$$

where  $\omega$  is a material parameter,  $g_{s_0}$  is the initial saturation hardness, and  $\dot{\gamma}$  is the total slip rate over all slip systems,

$$\dot{\gamma} = \sum_{\alpha} |\dot{\gamma}^{\alpha}|. \quad (17)$$

Note that setting  $\omega$  to zero results in a fixed value for the saturation hardness,  $g_s = g_{s_0}$ . This precipitation hardening law is well-suited to model strong self-hardening, such as results from the Orowan looping mechanism.

The microscale physics is implemented in the open-source Albany multiphysics simulation software. A fully implicit integration scheme is used to solve the governing equations for the constitutive update. For each integration point, we solve a fully-coupled system of 24 equations for both the slip  $\gamma^{\alpha}$  and the hardening  $g^{\alpha}$  at step  $n + 1$ .

## Microscale simulations


In this study, the crystal plasticity parameter values given in Table 3 were chosen to reflect a stainless steel alloy. The model was applied to simulate the behavior of the polycrystalline microstructure described in Section 2.1. Employing the microstructure illustrated in Figure 3, uniaxial stress simulations were conducted with symmetry boundary conditions applied to three planes of the body. A constant displacement rate resulting in an engineering strain rate of  $0.001 s^{-1}$  was applied along the long axis of the body. An engineering strain of 0.05 was achieved over ~200 time steps. The resulting mechanical behavior of the polycrystal is shown in Figure 5, where the load-displacement responses of the ten orientation realizations are plotted.

**Table 3.** Crystal plasticity parameters used in this study to model a representative stainless steel.

Elasticity			Flow		Hardening			
$C_{11}$ (GPa)	$C_{12}$ (GPa)	$C_{44}$ (GPa)	$m$	$\dot{\gamma}_0$ ( $s^{-1}$ )	$\dot{g}_0$ (MPa)	$g_{s_0}$ (MPa)	$g_0$ (MPa)	$\omega$
204.6	137.7	126.2	0.05	1	$2 \times 10^4$	202	90	0.01

## Macroscale model

The results of the microscale simulations referenced in the previous section were used to fit parameters for a standard  $J_2$  rate-independent plasticity model. While the crystal plasticity model used



**Figure 5.** Load-displacement response of ten realizations of crystal orientations.

in the microscale simulations is a rate-dependent plasticity model, the concurrent multiscale simulation is performed at a relatively low far-field strain rate (0.001/s) and samples the macroscale model over limited strain rates. In efforts to generalize the results presented here, future work is needed to address the effects of disagreement between macroscale and microscale material models regarding strain rate, plastic anisotropy, and multiphysics. For this study, the analysis focuses exclusively on the behavior of the microscale model within the concurrent multiscale framework.

In an attempt to reduce the data in a manner consistent with experimental data, engineering quantities were converted to true stress and true strain through the volume preservation assumption. To avoid convolving structural softening with the monotonic hardening exhibited by the model for homogeneous deformation, the microscale data were terminated for the fitting process at a strain of 0.04, which is prior to the onset of necking. In conjunction with the  $J_2$  yield function, the nonlinear hardening model of [77],


$$\sigma_y = \sigma_0 + H\epsilon_p + S(1 - e^{-\alpha\epsilon_p}) \quad (18)$$

was employed, with a yield surface partitioned through a yield stress  $\sigma_0$ , linear hardening term  $H$ , and a saturation hardening term governed by  $S$  and  $\alpha$ . The yield stress  $\sigma_0$  is chosen to be the departure from linearity at ~144 MPa. Small elastic strains were extracted in uniaxial stress through the assumption of isotropy with a Young's Modulus of 196 GPa. Fortunately, the parameterization

of the  $J_2$  model is sufficiently straightforward to permit rapid sensitivity studies and a manual fit that roughly minimizes the error by capturing mean behavior with continued hardening for strains beyond the simulated necking process. The resulting parameters are noted in Table 4. The reduced stress-strain curves and the proposed  $J_2$  fit are illustrated in Figure 6.

**Table 4.**  $J_2$  plasticity parameters fit to the reduced crystal plasticity simulations.

Elasticity		Hardening			
$E$ (GPa)	$\nu$	$\sigma_0$ (MPa)	$H$ (MPa)	$S$ (MPa)	$\alpha$
195.0	0.3	144	300	170	190



**Figure 6.** Reduced true stress-strain curves and the  $J_2$  fit of ten realizations of crystal orientations.

## 2.2 Concurrent multiscale modeling with the Schwarz Alternating Method

In order to concurrently model a domain combining the micro- and macroscale model behavior detailed above, the Schwarz alternating method was applied. The Schwarz alternating method was introduced by Hermann Schwarz in 1870 [76] to prove the existence of harmonic functions in irregular domains by expressing them as the union of domains of regular shape, for example



rectangles and circles. Thus understood, the method relies on knowledge of analytical solutions on simple domains, but the same framework can be applied to couple numerical solutions on multiple overlapping subdomains to achieve a solution on a combined domain. Recently, the method has been employed to couple atomistic and continuum scales in [29, 63, 65, 83], to couple distinct physics in different subdomains such as fluid-structure interaction [24] and computational fluid dynamics with aero-acoustics [15, 73]. Recently, the Schwarz alternating method has been demonstrated for use in concurrent multiscale continuum-to-continuum coupling of an arbitrary number of subdomains for finite-deformation solid mechanics, along with proof of a theoretical guarantee of convergence under certain restrictions that are typically met in such analyses [51]. A detailed description of the method is found in [51], but a brief summary is given below.

The Schwarz alternating method proceeds to solve a BVP with associated energy functional  $\Phi[\varphi]$  on a domain  $\Omega$  by solving a sequence of BVPs with associated energy functionals  $\Phi_i[\varphi]$  on a set of  $N$  subdomains  $\Omega_{1\dots N}$  with  $\cup_{i=1}^N \Omega_i = \Omega$ . The induced ordering is arbitrary except that there must be an overlap region between consecutive subdomains; that is,  $\Omega_i \cap \Omega_{i+1} \neq \emptyset$  for  $i \in \{1, N-1\}$  and  $\Omega_N \cap \Omega_1 \neq \emptyset$ . The boundary conditions on the subset  $\partial_\phi \Omega$  of the boundary of the combined domain are augmented with Schwarz boundary conditions on the subsets  $\Gamma_i$  of the boundaries of the subdomains  $\Omega_i$ , as follows. The values of the primary solution variables on domains  $2 \dots N$  are typically initialized with the trivial solution, the solutions from these subdomains are interpolated onto  $\Gamma_1 = \partial\Omega_1 \cap \cup_{i=2}^N \Omega_i$ , and then the augmented BVP is solved on  $\Omega_1$ . The solution from  $\Omega_1$  is then interpolated on the boundary  $\Gamma_2$  of the next subdomain  $\Omega_2$ , and the BVP is solved therein. Subsequently, the interpolated solution in  $\Omega_2$  is transferred as a boundary condition to the next subdomain, and so on through the  $N$  subdomains. Finally, the solution is computed on the last subdomain  $\Omega_N$ , and the solution from this last subdomain is interpolated onto the boundary  $\Gamma_1$  of the first subdomain  $\Omega_1$ , and the cycle continues until convergence is achieved. In this study, a tolerance of  $10^{-6}$  applies to the norm of the global residual to indicate convergence. The iterative solution process is summarized in Algorithm 1.

---

**Algorithm 1** Schwarz Alternating Method

---

```

1:  $\varphi^{(0)} \leftarrow \text{id}_X$  in  $\Omega_{2\dots N}$  ▷ initialize to zero displacement or a better guess in  $\Omega_{2\dots N}$ 
2:  $n \leftarrow 1$ 
3: repeat ▷ Schwarz loop
4:    $i \leftarrow \text{mod}(n-1, N) + 1$  ▷ Cycle through  $N$  domains
5:    $\varphi^{(n)} \leftarrow \chi$  on  $\partial_\phi \Omega_i$  ▷ Dirichlet BC for  $\Omega_i$ 
6:    $\varphi^{(n)} \leftarrow P_{\Omega_j \rightarrow \Gamma_i}[\varphi^{(n-1)}]$  on  $\Gamma_i$  ▷ Schwarz BC for  $\Omega_i$ 
7:    $\varphi^{(n)} \leftarrow \arg \min_{\varphi \in \mathcal{S}_i} \Phi_i[\varphi]$  in  $\Omega_i$  ▷ solve on  $\Omega_i$  in the space of admissible functions  $\mathcal{S}_i$ 
8:    $n \leftarrow n + 1$ 
9: until converged


```

---

As proved in [51], this iterative procedure is guaranteed to converge linearly to the solution  $\arg \min_{\varphi \in \mathcal{S}} \Phi[\varphi]$  on the combined domain for the finite deformation solid mechanics problem. It is noted that the rate of this convergence increases with increasing overlap between the subdomains, so that the computational demands of computing constitutive response in the overlap regions should be balanced with the cost of recomputing equilibrium in order to optimize the over-

all cost.

In this study, the Schwarz method is applied to the domain shown in Figure 7, which comprises an ASTM standard tensile specimen scaled down so that the square gauge section has dimensions  $100\mu\text{m} \times 100\mu\text{m}$ . Each grip end of the specimen is modeled as a separate subdomain, and the




**Figure 7.** Subdomains of tensile specimen. The domain is split into two grip subdomains, each colored blue, and a microstructural subdomain at the center of the gauge section, colored by grain number. The  $50\mu\text{m}$  overlap regions can be seen at either end of the microstructural subdomain.

central part of the gauge section is a third subdomain. There is an overlap of  $50\mu\text{m}$  of the gauge subdomain with each of the grip subdomains, which is chosen to balance the computational costs of computing global equilibrium with the costs of computing the local constitutive response. The macroscale model detailed in Section 2.1 is employed in the grip subdomains, where the mesh has a characteristic element length on the order of  $10\mu\text{m}$ , and the microscale model discussed in Section 2.1 is employed in the gauge section subdomain, where the characteristic element length is  $2.5\mu\text{m}$ . It is noted that this difference in element size translates to a reduction in the condition number of the global tangent stiffness of 2-3 orders of magnitude in the Schwarz simulation compared with a solution on a monolithic domain with a graded mesh. The following section presents the results of applying this simulation setup to study the onset of necking localization.

### 3 Concurrent multiscale: Findings

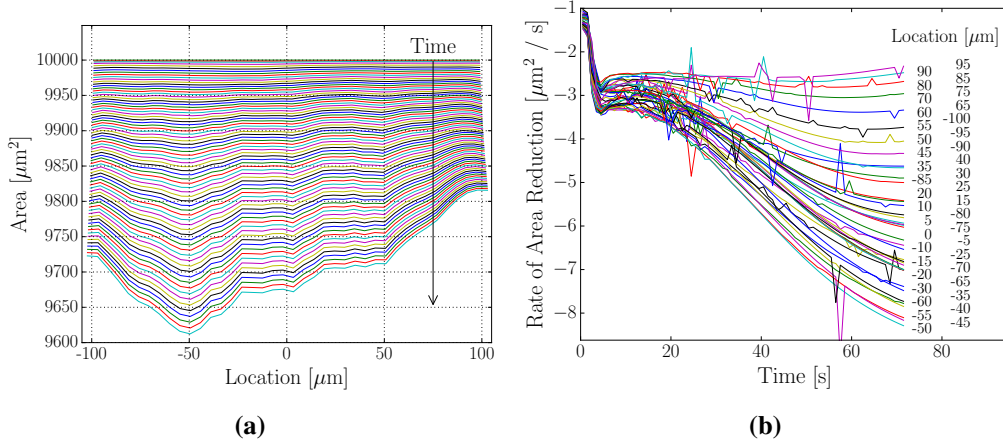
The multiscale simulation is conducted to apply a loading of approximately uniaxial stress to the gauge section of the tensile specimen. The inside surfaces of the holes in the grips are rigidly displaced at a rate of  $10^{-4}$  mm/s to achieve a constant engineering strain rate of approximately  $10^{-4}$  s $^{-1}$  as measured from the grip displacement. The simulation is run until this strain measure reaches approximately 0.5%. Due to the localization of the deformation in the gauge section, the strain measured from the ends of the microscale subdomain at the end of the simulation is close to 4%. Furthermore, as seen in Figure 8, the local strain approaches 20% where necking is incipient, representing more than an order of magnitude increase over the macroscopic strain level and five times the strain that would be observed with a typical experimental extensometer-based strain measurement. To provide a more quantitative measure of localization, the cross-sectional



**Figure 8.** Contour of  $F_{11}$ —the normal component of the deformation gradient in the loading direction. Localization is evident in the regions with red coloring.


area of the microscale domain is measured at each load step during the simulation. These measurements are performed by computing the convex hull of nodes comprising  $5\mu\text{m}$  slices of the domain and measuring the area in each slice. In Figure 9(a), the change in this cross-sectional area is plotted as a function of position along the axial direction. Near the position of  $-50\mu\text{m}$ , the area is seen to be reducing more rapidly than in the rest of the gauge section, indicating the onset of necking in that region. This trend is also evident in Figure 9(b), where the rate of area reduction is shown to be consistently greater in the vicinity of the  $-50\mu\text{m}$  position throughout the simulation.

The importance of employing concurrent multiscale coupling to transfer structural loads to the microstructure becomes clear when comparing these results to traditional methods of imparting deformation through "fixed" or "symmetry" boundary conditions. The corresponding reduction in area is plotted for the three sets of boundary conditions in Figure 10. In the figure, the label *Symmetry BC* indicates the application of a symmetry boundary condition, where homogeneous



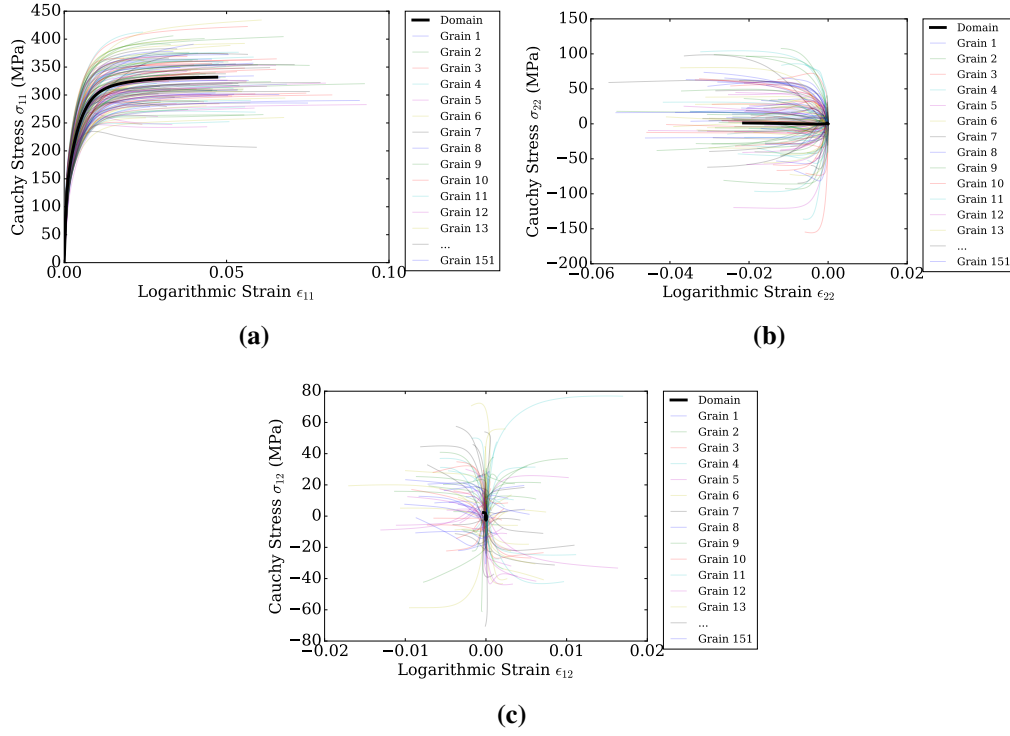
**Figure 9.** (a) Reduction in cross-sectional area over time. Localization is evident in the vicinity of  $-50\mu\text{m}$ , where the area is being preferentially reduced. (b) Rates of area reduction as a function of time. The signature of localization is seen in locations around  $-50\mu\text{m}$ , as in (a).

x-, y-, and z-displacement boundary conditions are applied to the negative x-, y-, and z-faces of the microstructure, respectively. In the *Fixed BC* case, homogeneous y- and z-displacement conditions are imposed at both ends. In both of these cases, the x-displacement is subject to a spatially uniform time-dependent boundary condition that maintains parallel and planar opposing x-faces. The same set of grain orientations is applied in each simulation. The data in Figure 10 is extracted from the three simulations when the minimum cross-sectional area is roughly  $9600\mu\text{m}^2$ . It is clear from these results that the fixed boundary conditions are inappropriate for modeling necking due to the overconstraint in the plane of the cross-section at the ends of the specimen. This is certainly not surprising, but it bears mention, as this is one of the most commonly applied boundary conditions in studying microstructural response. The symmetry and Schwarz boundary conditions produce more realistic profiles; however, in both cases, localization is most pronounced in the vicinity of the applied boundary conditions. For the symmetry boundary conditions, this is the end of the specimen, and for the Schwarz boundary conditions, this is the end of the overlap region, which is located at  $-50\mu\text{m}$ . For the symmetry boundary conditions, there is overconstraint of the x-displacement at the end of the specimen where necking is incipient. Whether this drives the location of the localization is not clear from a single simulation. For the Schwarz case, more study is required to determine whether this is a coincidence, or whether there is additional stress concentration at the boundary due to overconstraint owing either to the larger element size or the dissimilar material model employed in the grip section. If this is the case, mesh refinement can be used to eliminate the former problem, and material model iteration could eliminate the latter problem. A clear takeaway from these results is that careful choice of boundary conditions is critical to obtain quantitative information about necking.



**Figure 10.** Cross-sectional area profiles for simulations employing symmetry, fixed, and Schwarz boundary conditions.

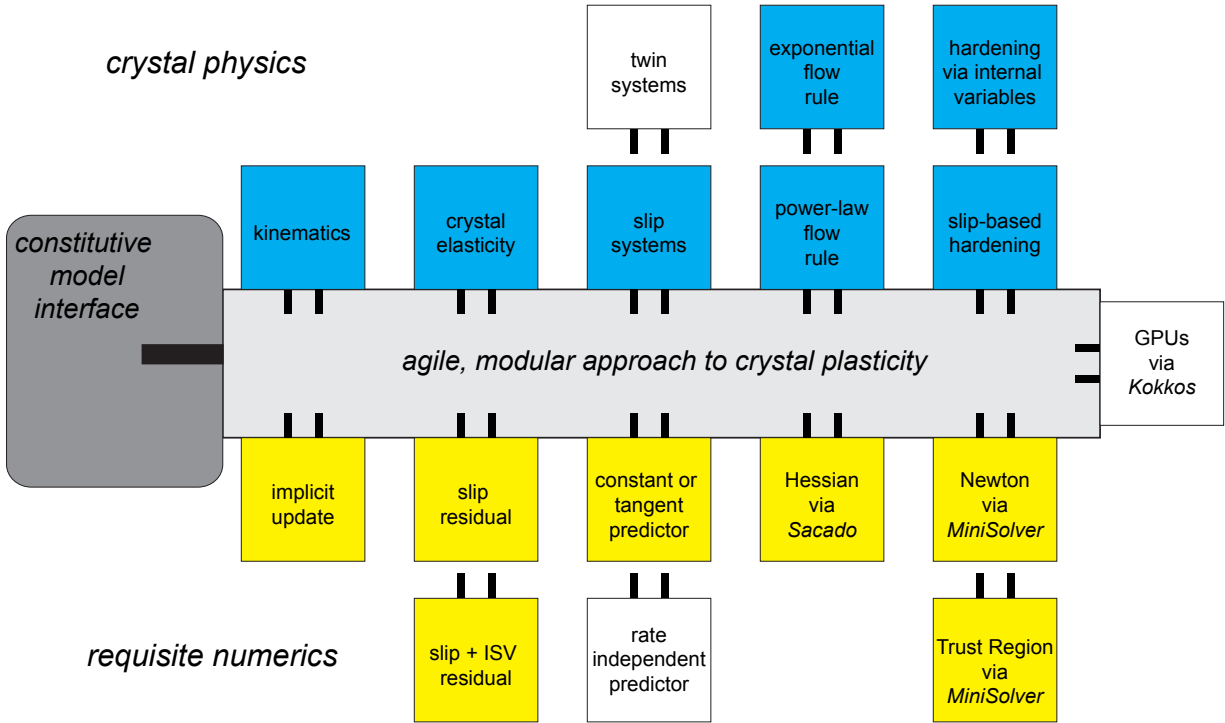
In Figure 11, the stress-strain response of the gauge section is plotted. The whole-domain average stress and strain, plotted with the solid lines, show the expected trends, as follows. The normal stress  $\sigma_{11}$ , shown in Figure 11(a), increases monotonically with the strain  $\epsilon_{11}$  in a manner typical of elasto-viscoplasticity with hardening. The transverse stress  $\sigma_{22}$ , shown in Figure 11(b), is approximately zero throughout the simulation while the transverse strain  $\epsilon_{22}$  decreases monotonically due to the expected Poisson effect, indicating a uniaxial stress condition. Lastly, as shown in Figure 11(c), the shear stress  $\sigma_{12}$  and shear strain  $\epsilon_{12}$  are nearly zero throughout the simulation, indicating the expected decoupling of the shear and normal components of the stress and strain. Also plotted in Figure 11 are the average stresses and strains associated with each grain in the microstructure. These quantities exhibit much more interesting behavior, with complicated anisotropy, non-uniaxial stress states, non-monotonic behavior, and coupling of normal and shear deformation. One striking feature is that the magnitude of the local transverse normal and shear stresses approach 40% and 20% of the average axial normal stress, respectively. This is a strong indication that failure models that are sensitive to triaxial or shear loading require local knowledge of the stress state to function effectively, and that the current framework will make this information available to such models in a way that the macroscale phenomenological model cannot.



**Figure 11.** Stress-strain response of microstructural domain: (a) Normal, along tensile axis; (b) Normal, perpendicular to tensile axis; (c) Shear in transverse plane. The solid lines with closed symbols correspond to the average response of the entire microstructural domain. The dashed lines correspond to the grain-averaged quantities.

## 4 Progress in crystal plasticity

During the inception of this program, we initiated a refactor of the crystal plasticity model initiated by Reese Jones in the LCM research environment. As a team, we agreed that an implicit update of the slips and the internal variables governing the hardening was needed. We also sought to provide developers with multiple residuals that stemmed from crystal structure, elasticity, flow rules, and hardening laws. Through automatic differentiation and new packages for solution, we can accommodate the complexity inherent in discovery. A notional figure illustrating both the dominant crystal physics and the requisite numerics is illustrated in Figure 12. Instead of viewing the constitutive model as input physics that “plug” into our finite element code via an interface, we envision a modular model having multiple physics that can be interchanged to capture the dominant mechanisms for the crystal of interest.



**Figure 12.** A modular approach to crystal plasticity having the desired crystal physics and requisite numerics.

Resolving the complex and coupled physics requires a modular method for solution that incorporates multiple residuals, predictors, and solution methodologies. MiniSolver provides a framework for the solution of the local, highly nonlinear system of equations that arises from an implicit integration scheme of the constitutive model. MiniSolver is a component of the MiniTensor

library, which in turn is part of the Trilinos library for numerical computation.\* MiniSolver provides solution schemes such as Newton-Raphson, conjugate gradient, trust region, different types of line search, etc., in a manner that is independent of the implementation details of the constitutive model. Thus, it relieves the material model implementer from the tedious task of programming a solution scheme, and also provides the ability to easily change solution methods. In addition, the MiniTensor library provides the ability to write vectorial and tensorial expressions in a compact manner, much like it is done in textbooks and journal articles. This reduces the difference between the mathematical expression of a constitutive model and its representation in computer code, thus improving readability and reducing implementation errors. Other aspects of illustrated modularity are reflected in the extension to manycore architectures through Kokkos which is addressed in Section 4.5. We highlight progress in this section through small contributions in microscale physics, methods of solution, verification, maturation, and manycore applications.

## 4.1 Viscoplasticity via multiple flow rules and hardening laws

The physics of crystal plasticity is subject to phenomenology that varies with material, temperature, and strain rate. In the algorithms developed in this project, different approaches to modeling crystal plasticity are subsumed into a single model encompassing the combinatorial complexity arising from the modular design. The two main features of this implementation are a choice of flow rule and a choice of hardening laws, which can be chosen in any available combination via directives in an input file.

For the flow rule, the standard power law form of Equation (13) is one choice. A second option is the hybrid flow rule in Equation (22), which extends the power law flow rule to apply in the intermediate to high strain rate regime. A third option is the thermally activated flow rule,

$$\dot{\gamma}^\alpha = \dot{\gamma}_0^\alpha \exp \left\{ -\frac{F_0}{k_B T} \left\langle 1 - \left\langle \frac{\tau^\alpha - s_a^\alpha}{s_t^\alpha} \right\rangle^p \right\rangle^q \right\} \quad (19)$$

which is typically employed for materials with higher Peierls stress, such as body-centered cubic (BCC) metals.

For hardening laws, one choice is the saturation hardening law given by Equations (15), (16), (17). Additional choices are linear hardening with recovery,

$$\dot{g}^\alpha = \dot{g} = (H - R_d g) \dot{\gamma} \quad (20)$$

and a dislocation-density based formulation,

$$g^\alpha = A \mu b \sqrt{H^{\alpha\beta} \rho^\beta} \quad (21)$$

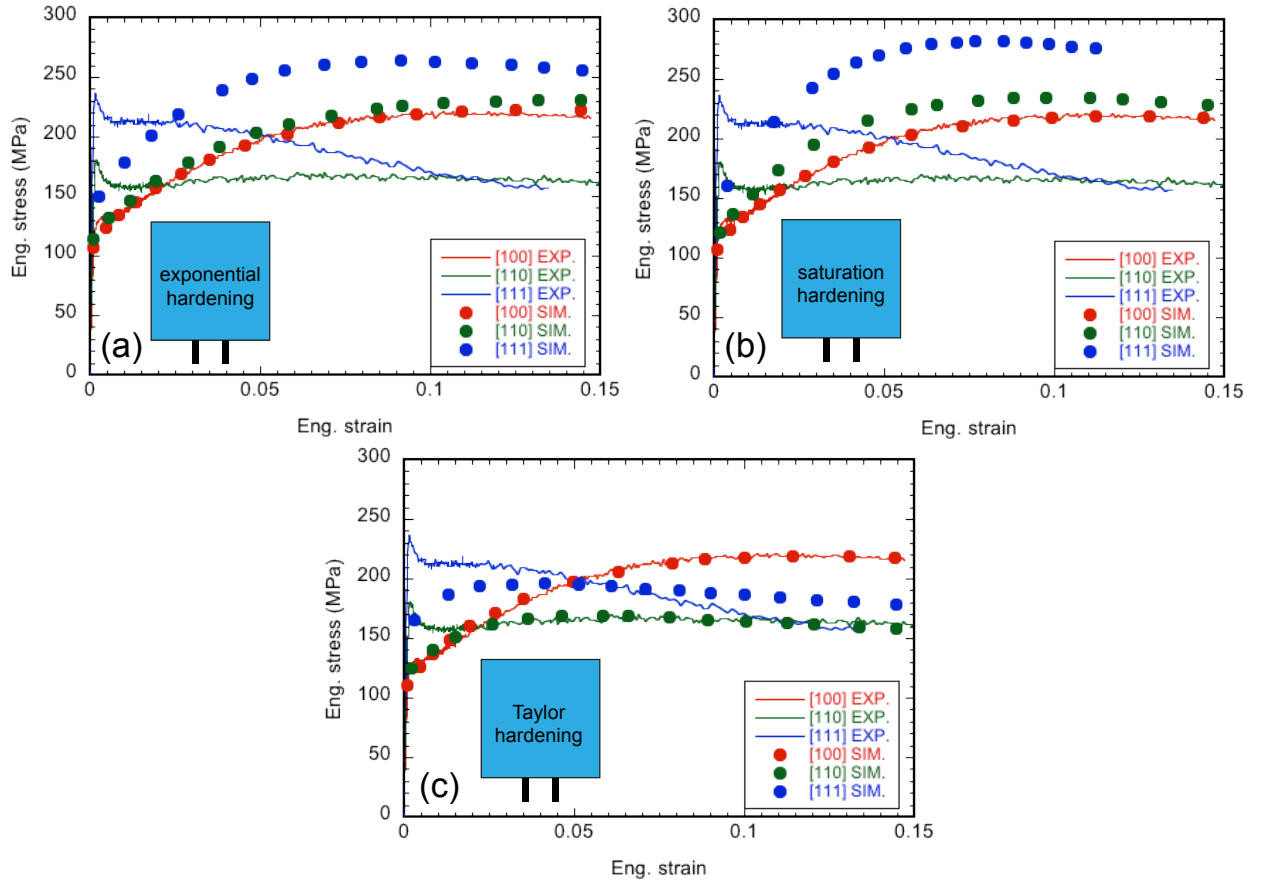
As an example of the importance of including this modularity, it was shown that for tantalum single crystal material behavior, the choice of hardening law can lead to a predictive model or a

---

\*Trilinos is available on GitHub: <https://github.com/trilinos>



prescriptive model calibration. This result is represented in Figure 13, where experimental [100] data is used to calibrate the hardening law parameters, and the model predicts behavior for [110] and [111] tests.



**Figure 13.** Comparing exponential hardening, saturation hardening, and Taylor hardening for tantalum single crystals of multiple orientations. Models were fit with experimental data in the [100] orientation and then employed to predict the response in the [110] and [111] orientations.

## 4.2 Incorporating drag-based flow with nonlinear root-finding

The traditional power-law relation shown in Equation (13) that dictates the flow rule exponentially increases when the resolved shear stress on a given system  $\alpha$  exceeds its corresponding hardening. As the sensitivity factor  $m$  decreases, the flow rule approaches the limiting case representing elastic perfectly plastic behavior, and the transition into the regime of unbounded slip rates becomes increasingly sharp [35]. Solving the nonlinear ODE numerically becomes a computationally arduous


ous task since any iterate of the slip that produces a sufficiently high resolved shear stress will lead to a high residual and Hessian matrix thus resulting in sub-optimal convergence rates, as shown for a single slip system in Figure 14(a). It is important to note that the high shear stresses produce slip rates that are non-physical, and are purely a result of the definition of the power-law, which is primarily accurate in the low slip-rate regime. Therefore, it is desirable to introduce a flow rule that is both physically admissible and contains bounded derivatives in the high stress regime. A flow rule that satisfies these requirements can be constructed by defining a continuous partitioning of the power-law dislocation motion with one that is viscous drag-dominated as follows

$$\dot{\gamma}^\alpha = \dot{\gamma}_0^\alpha \left( \frac{1}{\dot{\gamma}_{PL}^\alpha} + \frac{1}{\dot{\gamma}_{VD}^\alpha} \right)^{-1},$$

$$\dot{\gamma}_{PL}^\alpha = \left| \frac{\tau^\alpha}{g^\alpha} \right|^{\frac{1}{m}} \text{sgn}(\tau^\alpha), \quad (22)$$

$$\dot{\gamma}_{VD}^\alpha = \left( \frac{\tau^\alpha}{Dg^\alpha} \right),$$

where  $D$  is a user-specified drag parameter that is assumed to be constant. We are defining an effective slip rate  $\dot{\gamma}^\alpha$  through the micromechanical assumption of an additive process [11]. The flow rule shown above maintains the original power-law behavior when the resolved shear stress on a given slip system is low relative to the hardening. Once the shear stress increases, the drag term begins to dominate the flow rule, and eventually leads to a slip rate that grows linearly with increasing shear stress, as is shown for a single slip system in Figure 14(b).



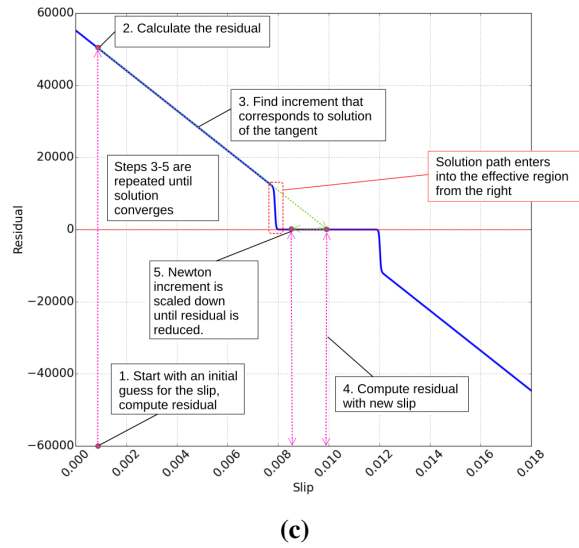
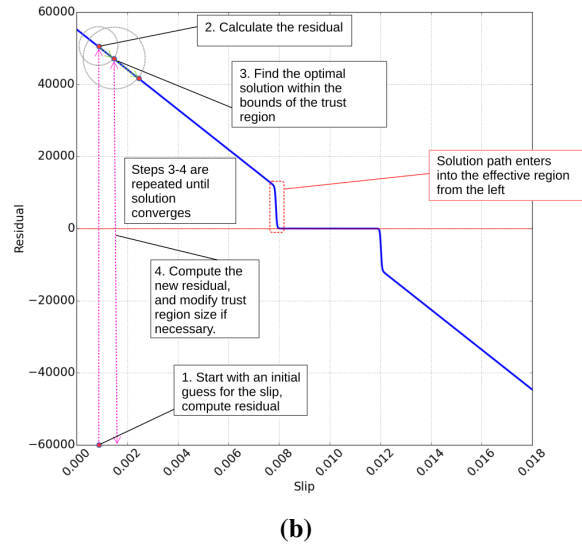
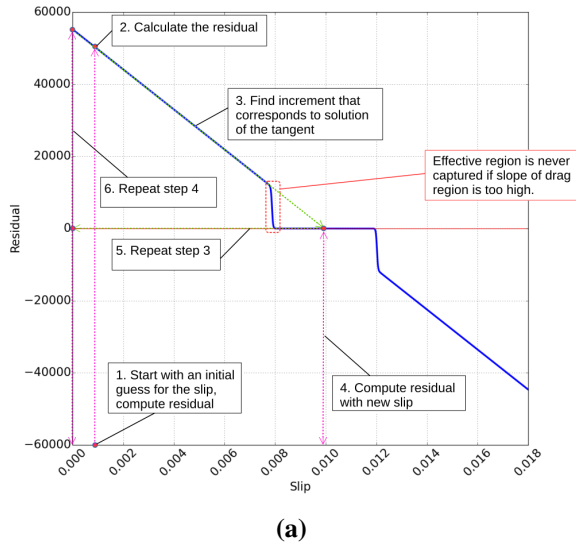
**Figure 14.** Residual profile for a single slip system using (a) the power-law flow rule, and (b) the power-law with drag flow rule.

Solving the above ODE for the slips at a given time step  $t_{n+1}$  poses new difficulties using


standard root-finding algorithms such as Newton’s method. Rapid variations in slope occur within the narrow region of slips where both the drag and the power-law terms are dominant. Moreover, the slope remains constant in the other regions where either the drag term or the power-law term dictate the flow rule. Using Newton’s method to solve for the slips will therefore result in oscillations between the power-law and drag regimes unless the initial guess for the slips lies in the narrow region that contains a non-constant slope. As an alternative, other root-finding algorithms such as trust region or back-tracking line search can be used to constrain the space of admissible iterative increments [55]. The trust region algorithm attempts to obtain a solution using a specified bandwidth that is centered around the current iterate which can grow or shrink depending on certain parameters that indicate the sensitivity of the residual to small perturbations in the slips. This algorithm can therefore prevent iterations from oscillating between values that are significantly different. The back-tracking line search similarly constrains the space where the solution for the slips is sought. This methodology uses an iterative increment obtained using Newton’s method, and applies a uniform scaling factor to ensure that the subsequent residual is decreased by a sufficient amount [7]. Both root-finding algorithms are generally robust for single slip systems, as is shown in Figure 15. The cumulative time based on the adaptive time-stepping as well as the typical number of material solver iterations per global step is shown in Figure 16 for a simple cube with an imposed uniaxial strain-rate. These figures indicate that using the drag-based flow rule both reduces the number of iterations required to obtain a solution, and permit a larger average time step throughout the simulation.

The multi-slip scenario significantly changes the nature of the ODE for the slips. Figure 17 plots the natural logarithm of the absolute value of the residual for a given slip system of a material that contains 12 slip systems. Each point represents a state of 12 random slips. In addition, larger circles indicate a more prominent drag on the given slip system. This figure clearly illustrates that a fixed value of slip on the current iterate can result in a widely varying value of the residual, since it depends on the state of all 12 systems.


A model problem used to test the robustness and accuracy of the numerics is the Rubik’s cube block shown in Figure 18 that consists of 27 equal-sized blocks, each with random crystal orientations and a 12 slip-system FCC lattice structure. Considering a representative material point at the initial simulation step and a sufficiently small time step to ensure convergence, Figure 19 relates the iterative process of each component of the slips and residual to the 2-norm of the 12-component residual. Figure 19(a) illustrates a monotonic convergence using the power-law flow rule without drag. It is evident that with a sufficiently small time step, each iteration preserves the dominant slip systems, without generating finite slips in other components that are not initially active. Adding the drag term as shown in Figure 19(b) initially leads to a different set of active slips which is corrected mid-way through the iterative process thus leading to a sudden increase in the total residual. By limiting the rate at which the slips can grow at high strain-rates, the drag-based flow rule enables the activation of certain slip systems that are not naturally prominent in the power-law flow rule. In order to mitigate this effect, it is desirable to incorporate constraints into the root-finding algorithm that prevent the iterative solution from diverging or oscillating. In contrast to the single-slip example, the standard criterion for the line search that accepts an increment based on the desired reduction in the residual does not guarantee monotonic convergence, as is highlighted in Figure 19(b). Another reduction criterion can be selected to produce an increment



**Figure 15.** The incremental solution path for the power-law with drag residual of a single slip system using (a) Newton's method, (b) trust region, and (c) back-tracking line search.




**Figure 16.** Comparison of (a) the elapsed model time for a given simulation step, and (b) the total material model iterations per time step using Newton's method with and without drag, trust region, and back-tracking line search.



**Figure 17.** Profile of a single component of the power-law with drag residual using random combinations of slips for a 12 slip system material point.

that only contains slips that exhibit a finite power-law behavior, and is thus relatively far from the drag-dominated regime. This requirement prevents the Hessian matrix from containing individual components with significantly varying tangents. However, this criterion does not guarantee convergence, even within the narrow range of power-law dominated slips. Other possibilities that can be explored in the future include variants of the trust-region and line search algorithms, as well as other constrained optimization techniques such as the barrier method.




**Figure 18.** The Rubik's cube block setup.

### 4.3 Verification through quad-slip


The computational infrastructure employed in our research environment was verified through numerous regression tests that focused on particular features of the crystal plasticity implicit update. Because our implementation permits the specification of a single slip system, we first developed analytical/computational approaches to verification through single slip under uniaxial strain in *Mathematica* [84]. We verified the evolution of slip to 15 significant figures. Although single-slip is indeed helpful, we seek additional cases which permit specification of all 12 FCC slip systems with idealized loadings that activate a subset of the slip systems.

In this section, we consider a particular crystal orientation under uniaxial stress. We will show that the resulting slip is relegated to four slip systems and henceforth refer to this particular case as quad-slip. For all intents and purposes, the case reflects a single degree of freedom for solution. The magnitude of slip on each of the four slip systems is identical. Moreover, both  $\mathbf{F}_e$  and  $\mathbf{F}_p$  have a diagonal character. Given an imposed stress, we can solve for  $\mathbf{F}_e$  and essentially post-process slip with exponential updates for  $\mathbf{F}_p$ . To illustrate the construction and utility of this approach, we first define a rotated basis

$$\mathbf{e}_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}, \mathbf{e}_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}, \mathbf{e}_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \quad (23)$$



(a)



(b)

**Figure 19.** The iterative solution and residual path for a representative material point in the Rubik's cube problem using (a) power-law flow rule and Newton's method, and (b) power-law with drag flow rule and Newton's method.

In this rotated basis, the 6x6 matrix of the components of the elasticity tensor is

$$[\mathbb{L}] = \begin{bmatrix} \frac{1}{2}(C_{11} + C_{12} + 2C_{44}) & \frac{1}{2}(C_{11} + C_{12} - 2C_{44}) & C_{12} & 0 & 0 & 0 \\ \frac{1}{2}(C_{11} + C_{12} - 2C_{44}) & \frac{1}{2}(C_{11} + C_{12} + 2C_{44}) & C_{12} & 0 & 0 & 0 \\ C_{12} & C_{12} & C_{11} & 0 & 0 & 0 \\ 0 & 0 & 0 & C_{11} - C_{12} & 0 & 0 \\ 0 & 0 & 0 & 0 & 2C_{44} & 0 \\ 0 & 0 & 0 & 0 & 0 & 2C_{44} \end{bmatrix} \quad (24)$$

A uniaxial stress state is imposed; that is,

$$\boldsymbol{\sigma} = \begin{bmatrix} \sigma_{11} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad (25)$$

For the initial elastic step, this stress state is obtained when applying a velocity gradient of the form

$$\mathbf{L} = \mathbf{L}_e = \dot{\epsilon} \begin{bmatrix} 1 & 0 & 0 \\ 0 & -\nu_1 & 0 \\ 0 & 0 & -\nu_2 \end{bmatrix} \quad (26)$$

over a time increment  $\Delta t$ . Here,  $\dot{\epsilon}$  is the applied strain rate, and  $\nu_1$  and  $\nu_2$  are the appropriate Poisson's ratios. Assuming an additive decomposition of the velocity gradient into elastic and plastic parts, as in Equation (7), leads, via Equation (8) to the differential equation,

$$\dot{\mathbf{F}}_e = \mathbf{L}_e \cdot \mathbf{F}_e \quad (27)$$

The approximate update for  $\mathbf{F}_e$  is then

$$(\mathbf{F}_e)_{n+1} = \exp(\mathbf{L}_e \Delta t) \cdot (\mathbf{F}_e)_n \quad (28)$$

Assuming that  $(\mathbf{F}_e)_0 = \mathbf{I}$ ,  $\mathbf{F}_e$  will retain a diagonal form for its component matrix during the elastic steps; that is,

$$\mathbf{F}_e = \begin{bmatrix} (F_e)_{11} & 0 & 0 \\ 0 & (F_e)_{22} & 0 \\ 0 & 0 & (F_e)_{33} \end{bmatrix} \quad (29)$$

For  $\boldsymbol{\sigma}$  as in Equation (25),

$$\tau^\alpha = (F_e)_{11} (F_e)_{22} (F_e)_{33} \sigma_{11} (\mathbf{P}^\alpha)_{11} \quad (30)$$

For a diagonal  $\mathbf{F}_e$ ,  $(\mathbf{P}^\alpha)_{11} = (\mathbf{P}_0^\alpha)_{11}$ , and

$$\tau^\alpha = (F_e)_{11} (F_e)_{22} (F_e)_{33} \sigma_{11} (\mathbf{P}_0^\alpha)_{11} \quad (31)$$

The bases in Equation (23) are chosen such that  $(\mathbf{P}_0^\alpha)_{11}$  is nonzero only for slip systems  $\{7, 8, 10, 11\}$  in Table 2. The matrices of components in the rotated basis for the active slip systems are

$$\mathbf{P}_0^5 = \begin{bmatrix} -\frac{1}{\sqrt{6}} & 0 & \frac{1}{\sqrt{12}} \\ \frac{1}{\sqrt{6}} & 0 & -\frac{1}{\sqrt{12}} \\ -\frac{1}{\sqrt{3}} & 0 & \frac{1}{\sqrt{6}} \end{bmatrix} \quad \mathbf{P}_0^6 = \begin{bmatrix} \frac{1}{\sqrt{6}} & 0 & -\frac{1}{\sqrt{12}} \\ -\frac{1}{\sqrt{6}} & 0 & \frac{1}{\sqrt{12}} \\ \frac{1}{\sqrt{3}} & 0 & -\frac{1}{\sqrt{6}} \end{bmatrix} \quad \mathbf{P}_0^{11} = \begin{bmatrix} -\frac{1}{\sqrt{6}} & 0 & -\frac{1}{\sqrt{12}} \\ -\frac{1}{\sqrt{6}} & 0 & -\frac{1}{\sqrt{12}} \\ \frac{1}{\sqrt{3}} & 0 & \frac{1}{\sqrt{6}} \end{bmatrix} \quad \mathbf{P}_0^{12} = \begin{bmatrix} \frac{1}{\sqrt{6}} & 0 & \frac{1}{\sqrt{12}} \\ -\frac{1}{\sqrt{6}} & 0 & -\frac{1}{\sqrt{12}} \\ -\frac{1}{\sqrt{3}} & 0 & -\frac{1}{\sqrt{6}} \end{bmatrix} \quad (32)$$



From Equations (30) and (32),

$$\tau^6 = \tau^{12} = -\tau^5 = -\tau^{11} = \tau = \frac{1}{\sqrt{6}} (F_e)_{11} (F_e)_{22} (F_e)_{33} \sigma_{11} \quad (33)$$

Assuming that the flow rule  $\dot{\gamma}^\alpha (\tau^\alpha)$  is the same for each system,

$$\dot{\gamma}^6 = \dot{\gamma}^{12} = -\dot{\gamma}^5 = -\dot{\gamma}^{11} = \dot{\gamma} \quad (34)$$

This would require, for example, equivalent hardening on each system,

$$\dot{g}^6 = \dot{g}^{12} = \dot{g}^5 = \dot{g}^{11} = \dot{g} \quad (35)$$

If the slip resistance is initially the same on each system, and the evolution of the resistance is a function of the absolute value of the slip rate on the system, as is common, the flow rule is expected to remain the same for each system. Given the relationships in Equation (34), and the definition in Equation (9),

$$\mathbf{L}_p^* = 2\sqrt{\frac{2}{3}}\dot{\gamma} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \end{bmatrix} \quad (36)$$

Thus, given an applied velocity gradient with a diagonal component matrix in the rotated basis, the elastic and plastic parts of the velocity gradient will retain diagonal component matrices, and so too will  $\mathbf{F}_e$  and  $\mathbf{F}_p$ . Now, the Cauchy stress  $\boldsymbol{\sigma}$  is given in terms of the second Piola-Kirchhoff stress  $\mathbf{S}$  in the intermediate configuration by

$$\boldsymbol{\sigma} = \frac{1}{\det(\mathbf{F}_e)} \mathbf{F}_e \cdot \mathbf{S} \cdot \mathbf{F}_e^T \quad (37)$$

With  $\mathbf{F}_e$  diagonal, and only  $\sigma_{11}$  nonzero as in Equation (25), attention is restricted to

$$\sigma_{11} = \frac{1}{(F_e)_{11} (F_e)_{22} (F_e)_{33}} S_{11} \quad (38)$$

Substituting from Equation (38) into Equation (33),

$$\tau = \frac{1}{\sqrt{6}} S_{11} \quad (39)$$

The elasticity relation for the stress  $\mathbf{S}$  is written in terms of the elastic Lagrangian strain  $\mathbf{E}_e$  as

$$\mathbf{S} = \mathbb{L} : \mathbf{E}_e \quad (40)$$

For  $S_{11}$ , this implies

$$S_{11} = L_{11kl} (\mathbf{E}_e)_{kl} = \frac{1}{2} L_{11kl} [(\mathbf{F}_e)_{mk} (\mathbf{F}_e)_{ml} - \delta_{kl}] \quad (41)$$

For  $\mathbb{L}$  as in Equation (24), this implies

$$S_{11} = \frac{1}{2} \{ L_{1111} [(\mathbf{F}_e)_{m1} (\mathbf{F}_e)_{m1} - 1] + L_{1122} [(\mathbf{F}_e)_{m2} (\mathbf{F}_e)_{m2} - 1] + L_{1133} [(\mathbf{F}_e)_{m3} (\mathbf{F}_e)_{m3} - 1] \} \quad (42)$$

Given that  $\mathbf{F}_e$  is diagonal,

$$S_{11} = \frac{1}{2} \{ L_{1111} [(\mathbf{F}_e)_{11}^2 - 1] + L_{1122} [(\mathbf{F}_e)_{22}^2 - 1] + L_{1133} [(\mathbf{F}_e)_{33}^2 - 1] \} \quad (43)$$

Combining Equations (33), (38), and (43),

$$\tau = \frac{1}{\sqrt{24}} \{ L_{1111} [(\mathbf{F}_e)_{11}^2 - 1] + L_{1122} [(\mathbf{F}_e)_{22}^2 - 1] + L_{1133} [(\mathbf{F}_e)_{33}^2 - 1] \} \quad (44)$$

From  $S_{22} = 0$ ,

$$[(\mathbf{F}_e)_{22}^2 - 1] = -\frac{L_{2211}}{L_{2222}} [(\mathbf{F}_e)_{11}^2 - 1] - \frac{L_{2233}}{L_{2222}} [(\mathbf{F}_e)_{33}^2 - 1] \quad (45)$$

From  $S_{33} = 0$ ,

$$[(\mathbf{F}_e)_{33}^2 - 1] = -\frac{L_{3311}}{L_{3333}} [(\mathbf{F}_e)_{11}^2 - 1] - \frac{L_{3322}}{L_{3333}} [(\mathbf{F}_e)_{22}^2 - 1] \quad (46)$$

Combining Equations (45) and (46),

$$[(\mathbf{F}_e)_{22}^2 - 1] = \frac{L_{2233}L_{3311} - L_{3333}L_{2211}}{L_{2222}L_{3333} - L_{2233}L_{3322}} [(\mathbf{F}_e)_{11}^2 - 1] \quad (47)$$

Substituting back into Equation (46),

$$[(\mathbf{F}_e)_{33}^2 - 1] = \frac{L_{3322}L_{2211} - L_{3311}L_{2222}}{L_{2222}L_{3333} - L_{2233}L_{3322}} [(\mathbf{F}_e)_{11}^2 - 1] \quad (48)$$

Substituting from Equations (47) and (48) into Equation (44),

$$\begin{aligned} \tau &= \frac{1}{\sqrt{24}} \left\{ L_{1111} - \frac{L_{1122}L_{3333}L_{2211} - 2L_{1133}L_{3322}L_{2211} + L_{1133}L_{3311}L_{2222}}{L_{2222}L_{3333} - L_{2233}L_{3322}} \right\} [(\mathbf{F}_e)_{11}^2 - 1] \\ &= \tau_0 [(\mathbf{F}_e)_{11}^2 - 1] \end{aligned} \quad (49)$$

Now,

$$\left( \dot{\mathbf{F}}_e \right)_{11} = (\mathbf{L}_e)_{11} (\mathbf{F}_e)_{11} = [(\mathbf{L})_{11} - (\mathbf{L}_p)_{11}] (\mathbf{F}_e)_{11} = \left[ \dot{\epsilon} - 2\sqrt{\frac{2}{3}}\dot{\gamma} \right] (\mathbf{F}_e)_{11} \quad (50)$$

Taking  $\dot{\gamma}$  as in Equation (13),

$$\left( \dot{\mathbf{F}}_e \right)_{11} = \left[ \dot{\epsilon} - 2\sqrt{\frac{2}{3}}\dot{\gamma}_0 \left( \frac{\tau}{g} \right)^{1/m} \right] (\mathbf{F}_e)_{11} \quad (51)$$

Substituting from Equation (49),

$$\left( \dot{\mathbf{F}}_e \right)_{11} = \left[ \dot{\epsilon} - 2\sqrt{\frac{2}{3}}\dot{\gamma}_0 \left( \frac{\tau_0 [(\mathbf{F}_e)_{11}^2 - 1]}{g} \right)^{1/m} \right] (\mathbf{F}_e)_{11} \quad (52)$$

Taking  $\dot{g}^\alpha$  as in Equation (15), with the relationships in Equations (34) and (35), the evolution of the slip system resistance is governed by

$$\begin{aligned}\dot{g} &= \frac{4}{3}\dot{g}_0 \left( \frac{g_s - g}{g_s - g_0} \right) \dot{\gamma} \\ &= \frac{4}{3}\dot{g}_0 \dot{\gamma}_0 \left( \frac{g_s - g}{g_s - g_0} \right) \left( \frac{\tau_0 [(\mathbf{F}_e)_{11}^2 - 1]}{g} \right)^{1/m}\end{aligned}\quad (53)$$

The saturation hardness is

$$g_s = g_{s0} \left| 4 \left( \frac{\tau_0 [(\mathbf{F}_e)_{11}^2 - 1]}{g} \right)^{1/m} \right|^\omega \quad (54)$$

Substituting from Equation (54) into Equation (53),

$$\dot{g} = \frac{40}{3}\dot{g}_0 \dot{\gamma}_0 \left( \frac{g_{s0} \left| 4 \left( \frac{\tau_0 [(\mathbf{F}_e)_{11}^2 - 1]}{g} \right)^{1/m} \right|^\omega - g}{g_{s0} \left| 4 \left( \frac{\tau_0 [(\mathbf{F}_e)_{11}^2 - 1]}{g} \right)^{1/m} \right|^\omega - g_0} \right) \left( \frac{\tau_0 [(\mathbf{F}_e)_{11}^2 - 1]}{g} \right)^{1/m} \quad (55)$$

Equations (52) and (55) form a set of coupled, non-linear ordinary differential equations in  $(\mathbf{F}_e)_{11}$  and  $g$ . The saturation hardness is achieved when  $\dot{g} = 0$ . This is the case for  $(\mathbf{F}_e)_{11} > 0$  when  $g = g_s$ . That is, when

$$g = g_{s0} \left| 4 \left( \frac{\tau_0 [(\mathbf{F}_e)_{11}^2 - 1]}{g} \right)^{1/m} \right|^\omega \quad (56)$$

Solving Equation (56) for  $g$ ,

$$g = g_{s0}^{\frac{m}{m+\omega}} \left\{ 4^m \tau_0 [(\mathbf{F}_e)_{11}^2 - 1] \right\}^{\frac{\omega}{m+\omega}} \quad (57)$$

When the hardening has saturated, presumably the stress is constant, so that from Equation (52)


$$\begin{aligned}\dot{\epsilon} &= 2\sqrt{\frac{2}{3}}\dot{\gamma}_0 \left( \frac{\tau_0 [(\mathbf{F}_e)_{11}^2 - 1]}{g} \right)^{1/m} \\ &= 2\sqrt{\frac{2}{3}}\dot{\gamma}_0 \left( \frac{\tau_0 [(\mathbf{F}_e)_{11}^2 - 1]}{g_{s0}^{\frac{m}{m+\omega}} \left\{ 4^m \tau_0 [(\mathbf{F}_e)_{11}^2 - 1] \right\}^{\frac{\omega}{m+\omega}}} \right)^{1/m} \\ &= 2\sqrt{\frac{2}{3}} \frac{\dot{\gamma}_0}{4^{\frac{\omega}{m+\omega}} (g_{s0})^{\frac{1}{m+\omega}}} \left\{ \tau_0 [(\mathbf{F}_e)_{11}^2 - 1] \right\}^{\frac{1}{m+\omega}}\end{aligned}\quad (58)$$

Solving Equation (58) for  $(\mathbf{F}_e)_{11}$ ,

$$(\mathbf{F}_e)_{11} = \left\{ \frac{4^\omega g_{s0}}{\tau_0} \left( \sqrt{\frac{3}{8}} \frac{\dot{\epsilon}}{\dot{\gamma}_0} \right)^{m+\omega} + 1 \right\}^{1/2} \quad (59)$$

From Equations (39) and (49), then the saturated flow stress  $\bar{S}_{11}$  is

$$\bar{S}_{11} = \sqrt{6} 4^\omega g_{s0} \left( \sqrt{\frac{3}{8}} \frac{\dot{\epsilon}}{\dot{\gamma}_0} \right)^{m+\omega} \quad (60)$$



**Figure 20.** Constitutive response (solid line) for a strain rate of  $10^0 \text{ s}^{-1}$ . Model equations are (52) and (55), with parameters taken from Table 5. The dashed line is the saturated flow stress  $\bar{S}_{11}$  calculated from Equation (60).

**Table 5.** Crystal plasticity parameters employed for verification in uniaxial stress.

Elasticity			Flow		Hardening			
$C_{11}(GPa)$	$C_{12}(GPa)$	$C_{44}(GPa)$	$m$	$\dot{\gamma}_0$	$\dot{g}_0(MPa)$	$g_{s0}(MPa)$	$g_0(MPa)$	$\omega$
204.6	137.7	126.2	0.02	$10^{-3} \text{ s}^{-1}$	$2 \times 10^4$	202	90	0.01

#### 4.4 Porting to the LAME material library

The crystal plasticity constitutive model developed in this project was ported from the LCM code to the Library of Advanced Materials for Engineering (LAME), which is the material model library

used by the *Sierra/SolidMechanics* analysis code. Our strategy was to employ LCM as a vehicle for scientific exploration and algorithm development, and to transfer sufficiently mature technology to Sandia's ASC IC production codes where resources allowed. This approach is intended to maximize flexibility in the prototyping phase, while avoiding negative repercussions for the IC development teams.

Transferring technology that has been vetted in LCM to the IC codes is a primary mechanism for providing impact to the Sandia analyst community. In addition to providing direct analyst access to the constitutive model, the porting process also enables use of the constitutive model in conjunction with the vast *Sierra* toolset, for example in explicit transient dynamic simulations involving contact. In the case of the crystal plasticity constitutive model, source code was transferred directly from LCM to LAME. This was possible due to the LAME API, which provides a general interface for material models, and due to prior work by the *Sierra* and *Trilinos* teams to make *Trilinos* packages available within *Sierra/SM*.

The primary challenge in the porting process was adapting the LCM version of the material model to operate within the LAME and *Sierra/SM* infrastructures. The main tasks were:

- Implement *Sierra/SM* input deck parsing.
- Create conversion routines to load *Sierra/SM* input deck data into a `ParameterList` structure for use with the crystal plasticity model.
- Register proper state variables within the *Sierra/SM* framework.
- Implement routines for copying state data between *Sierra/SM* structures and crystal plasticity structures.
- Implement simplified routines for specifying lattice orientation within *Sierra/SM*
  - Provide a sensible set of predefined crystallographic structures (e.g., FCC, BCC).
  - Allow users to input orientations directly from the genesis mesh file via the variable initialization functionality in *Sierra/SM*.
- Strip all global-level automatic differentiation from the crystal plasticity model. Automatic differentiation is used in LAME within the constitutive model's state update routine, but is not passed between the material model and *Sierra/SM* as it is in LCM.
- Remove all use of LCM-specific data types from the crystal plasticity model.
- Write *Sierra/SM* tests that exercise the material model.

Care was taken in the porting process to ensure that both the LCM and *Sierra/SM* implementations imposed identical requirements on the genesis mesh file. That is, a mesh file created for use with LCM can be used directly in *Sierra/SM*, and vice versa. This is particularly important for discretizations of polycrystals because of the need to specify grain orientations and the potential for a large number of grains within a given mesh. Both LCM and *Sierra/SM* allow for the use of

polycrystal meshes in which each grain is tracked as a block, and in which the entire polycrystal is tracked as a single block. In the first case, the grain orientations may be prescribed for each element in the mesh file, or may be listed on a per-block basis in the input deck. In the second case, where the entire polycrystal is tracked as a single block, the grain orientations must be prescribed per element in the mesh file. Preprocessing tools were developed under this project to facilitate the creation of both types of polycrystalline meshes.

Several challenges related to the LCM and *Sierra* software development and testing environments impacted the porting process. The LCM source code was significantly modified to enable compatibility with the *Sierra* framework. A lasting consequence of these modifications is that it is not straightforward to maintain consistency between the two implementations, mainly because it is difficult to directly compare the source files. That is, the number of differences in the files that resulted from the porting process make it very difficult to identify differences that result from ongoing improvements to the model implemented in LCM. A second, high-level challenge results from differences in the build and test processes for LCM and *Sierra*. Specifically, the versions of *Trilinos* used by LCM and *Sierra* are generally not synchronized, which creates incompatibilities for newly-implemented *Trilinos* functionality. Additional challenges are tied to differences in compilers, compiler options, and test platforms for *Sierra* and LCM, which can result in code that compiles and passes the testing process in LCM, but either fails to compile or fails tests in *Sierra*, and vice versa.

## 4.5 Extension to manycore architectures

Porting large, complex scientific and engineering application codes to manycore architectures is complicated due to a variety of programming models, application programming interfaces (APIs), and performance requirements. A major challenge in utilizing heterogeneous resources is the diversity of devices on different computing platforms, which provide widely varying performance characteristics. A program or algorithm optimized for one architecture may not run as well on the next generation of processors or on a device from a different vendor. A program or algorithm optimized for GPU execution is often very different from one optimized for CPU execution. The relative performance of CPUs and GPUs also varies between computing platforms. On one machine, a specific portion of a computation may run best on the CPU, while on another machine it may run best on the GPU. In some cases, it is best to balance the workload between the CPU and the GPU; in other cases, it may be best to execute an algorithm on a device where it runs more slowly but closer to where its output is needed as to avoid expensive data transfer operations. This leads to a situation where an application code may have several implementations that are optimized for different architectures. In addition to the complexity of writing specialized codes, this approach leads to poor code maintainability: every new change to the code needs to be implemented in all versions.

Thus, the portability of the performance of a single code base is critical: parallel code needs to be executed correctly and remain performant on machines with different architectures, operating systems and software libraries. To this end, in this section we explore a new approach for achieving

performance portability in the LCM platform by using the Kokkos library.<sup>†</sup> Specifically, we target the constitutive model infrastructure in LCM for parallelization via Kokkos, as this can be done in a much shorter time than modifying other parts of the code.

We parallelized two materials models: a standard  $J_2$  plasticity model extended to the finite deformation regime, and the crystal plasticity model used throughout this report. A tailored material point simulator was written to determine the speedup in execution time with respect to the number of cores. Four sets of numerical experiments were conducted using the OpenMP parallel model available to Kokkos: 1000 and 1,000,000 instances of the  $J_2$  model, and 1000 and 10,000 instances of the crystal plasticity model. The results are shown in Figure 21. In the figure, the black dotted lines represent ideal speedup. The purple line represents the speedup for the total execution time, including parts of the code that have not been parallelized. The blue line represents the speedup for the numerical computation time, again including parts of the code that have not been parallelized. The red line, which is the more relevant here, represents the speedup of the constitutive model itself, which is the only part of the code that was parallelized for these experiments.

The results in Figure 21(a) show that with 1000 instances of the  $J_2$  model, the peak speedup is about 13 for 16 cores. Thereafter, adding cores to the computation leads to worse performance. By contrast, increasing the number of  $J_2$  instances to one million (Figure 21(b)) leads to increasing performance with increasing number of cores up to 32 cores, which is the maximum that was used. This suggests that the  $J_2$  model is relatively inexpensive in terms of computation, and thus, with as little as 1000 instances, there will be a point where adding more cores to a computation will starve them of work, leading to worse performance.


By contrast, we observe speedup for both 1000 and 10,000 instances of the much more computationally demanding crystal plasticity model, as shown in Figures 21(c) and 21(d). This suggests that to take full advantage of manycore architectures, starvation should be avoided. For relatively simple material models, this requires providing many instances of the material to each core to keep them busy. Complex material models like the crystal plasticity model appear to be sufficiently computationally intensive to keep cores busy even when using relatively few instances.

## 4.6 Robustness suite


During the course of this project, we attempted to focus on some aspects of test-driven development. We did devote resources to verification. We did investigate problems of increasing complexity for multiple numerical methods. We wanted, however, to begin to put some rigor behind particular studies so we could determine the robustness of particular input physics and numerical methods. To better examine these issues, we developed a python infrastructure to automatically populate and execute parametric studies for problems of increasing complexity. The kernel of the idea is expressed in Figure 22. Scripts incorporate parameters that govern multiple flow rules, hardening laws, numerical methods, discretizations, and time steps. A common theme in the analysis of polycrystals is the size of the implicit time step. We correlate a larger time step with increased

---


<sup>†</sup>Kokkos is available on GitHub: <https://github.com/kokkos>




(a)



(b)



(c)




(d)

**Figure 21.** Speedup for the OpenMP parallel implementation of the  $J_2$  and crystal plasticity models using Kokkos. (a) 1000 instances of the  $J_2$  model. (b) 1,000,000 instances of the  $J_2$  model. (c) 1000 instances of the crystal plasticity model. (d) 10,000 instances of the crystal plasticity model.



robustness and developed tools to mine log files for global iterations, convergence rates, and time steps.



**Figure 22.** We developed a robustness suite to quantify and monitor the performance of our algorithms for problems of increasing complexity.

Although the robustness suite is still in flux, one can access problems of increasing complexity through the SolidMechanicsExamples repository.

```
git clone https://code-source.sandia.gov/git/SolidMechanicsExamples
```

where the robustness suite is found in

```
SolidMechanicsExamples/LCM/CrystalPlasticity/robustnessTests
```

having an upper level python script `makeFiles.py`. Because of the combinatoric nature of these parametric studies, even relatively idealized problems can be computationally expensive. Hence, the true value of this suite will be realized when our cluster builds are reliable on *chama*, *skybridge*, *serrano*, and *ghost*. Although we have more to do regarding to making the robustness suite robust (the irony is quite painful), we believe that such a suite is necessary to both quantify and monitor the performance of our algorithms.

## 5 Advances in conformal microstructures

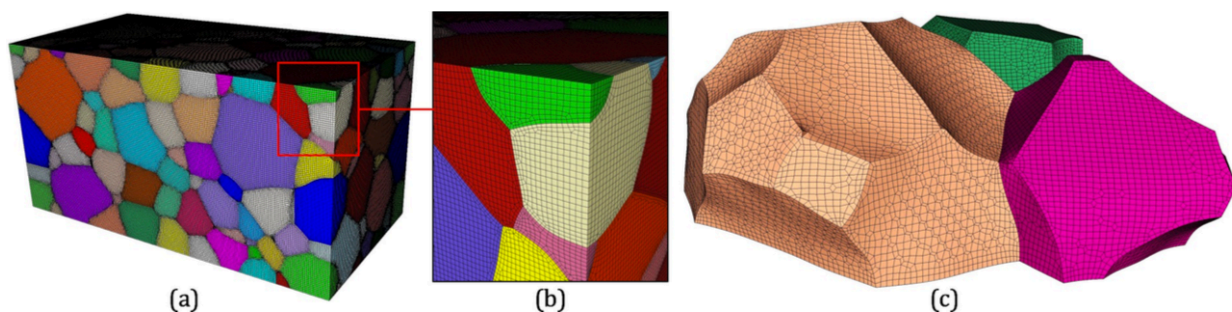
An accurate representation of microstructural features plays an important role in mesoscale computational models that explicitly resolve individual grains. The effect of detailed microstructural features becomes more pronounced when complex localized phenomena such as localization, void evolution, and eventual failure are considered. However, there exists some limitations to accurately reproduce these realistic 3D microstructures using finite element realizations. For example, the discretization of microstructures into finite elements generally produces voxelated interfaces between adjoining grains that create artificial stress/strain concentrations which can potentially trigger localization. The given element size convolutes the boundary representation with sufficient degrees of freedom to resolve the deformation process. Furthermore, synthetic microstructures are typically generated, where grain shapes are idealized or obtained from the output of a tiling (partitioning) protocol such as Voronoi tessellation. These microstructures, albeit easy to generate, are not informed by the energetics and thermodynamics of materials processes. In order to overcome these shortcomings, we present a modeling framework that incorporates physically-based realizations of polycrystalline aggregates into a finite element realization that conforms to interfaces. This technique utilizes Sandia National Laboratories' Cubit Sculpt meshing tool [58–61, 79]. It is shown that Cubit Sculpt provides a practical, robust and efficient tool to create interface conformal realizations using both hexahedral and tetrahedral elements. A general description of Sculpt algorithm is given by Owen [57] with recent extensions to microstructure [62]. In this section, we focus the application. We see to employ Sculpt to generate conformal grain boundaries with extensions to hexahedral and tetrahedral workflows.

### 5.1 Application to a hexahedral workflow

The interface conformal meshing technique has been successfully applied to various applications. In this section, we present some examples of computational microstructures and the generated meshes obtained from various techniques. Note that this meshing technique not only avoids voxelated representation of the microstructure, but also enables direct considerations of physically-based polycrystalline microstructures and thus a more realistic representation of microstructural features such as grain shapes, aspect ratios, grain boundaries and triple junctions. It is shown that the interface conformal FE mesh generated with the proposed methods is able to significantly reduce errors at grain boundaries and triple junctions in mechanical simulations when compared to an equivalent stair-step representation [43].

Figure 23 shows the finite element discretization of a polycrystalline microstructure generated by the Monte Carlo Potts grain growth model in SPPARKS. Here, a polycrystalline domain with 143 grains are represented by  $200 \times 100 \times 100$  voxels and their corresponding IDs from SPPARKS simulations. As shown in Figure 23 (c), Sculpt accurately reproduces smooth interfaces and junctions between grains.

Another class of computational methods used to generate physically-based microstructures is the phase-field (PF) grain growth models. The details of the PF grain growth model can be found




**Figure 23.** Polycrystalline microstructure generated from kinetic Monte Carlo grain growth simulations having 143 grains. (a) Full FEA mesh generated from  $200 \times 100 \times 100$  grid, (b) close-up of mesh, (c) view of three grains in FEA mesh.

in [1, 2]. In the PF model, grains are represented by the structural order parameters (OPs) and are evolved to satisfy the energy minimization. In this work, the total free energy is represented by the sum of bulk and interfacial energies, and the Allen-Cahn equation [6] is used for the dynamics. Figure 24 (a) shows 3D microstructure generated from the phase field (PF) grain growth simulation and the corresponding FE discretization [43]. It is shown that FE elements conform to the grain topology and ensure smoothness of the interfaces between adjoining grains. It is worth noting that since the interfaces in the phase field model are assumed to be isotropic in energy and mobility, balance of interfacial tensions at triple junctions, where three grains meet, yields angles of 120 degrees.


In addition to computational microstructure, Figure 25 shows the FE meshes generated from the experimental data created from electron backscatter diffraction (EBSD) measurement using a Scanning Electron Microscope (SEM). An EBSD technique characterizes microstructural features such as crystallography and phase of materials within a scanning grid. Crystallographic information from EBSD data is processed to have unique grain ids for each cell of a Cartesian grid that is used as input for the FE mesh generation procedure.

## 5.2 Moving to a tetrahedral workflow

Throughout the course of this effort, we were able to also invest in tetrahedral element technology. Gaps in Sculpt and the underlying MeshGems algorithm/implementation were highlighted and solved in FY16 and FY17 by Steve Owen and Corey Ernst. One of the improvements in the later part of FY16 was to move from a hexahedral representation of the entire domain to a triangle surface representation of the reconstructed grain boundaries. This enabled the implementation of smoothing algorithms specific to triangles. The geometric representation of the microstructure was broken into two components. The first component was a triangular surface mesh in an Exodus II format. The second component is a *s2g* file containing the requisite geometric information for

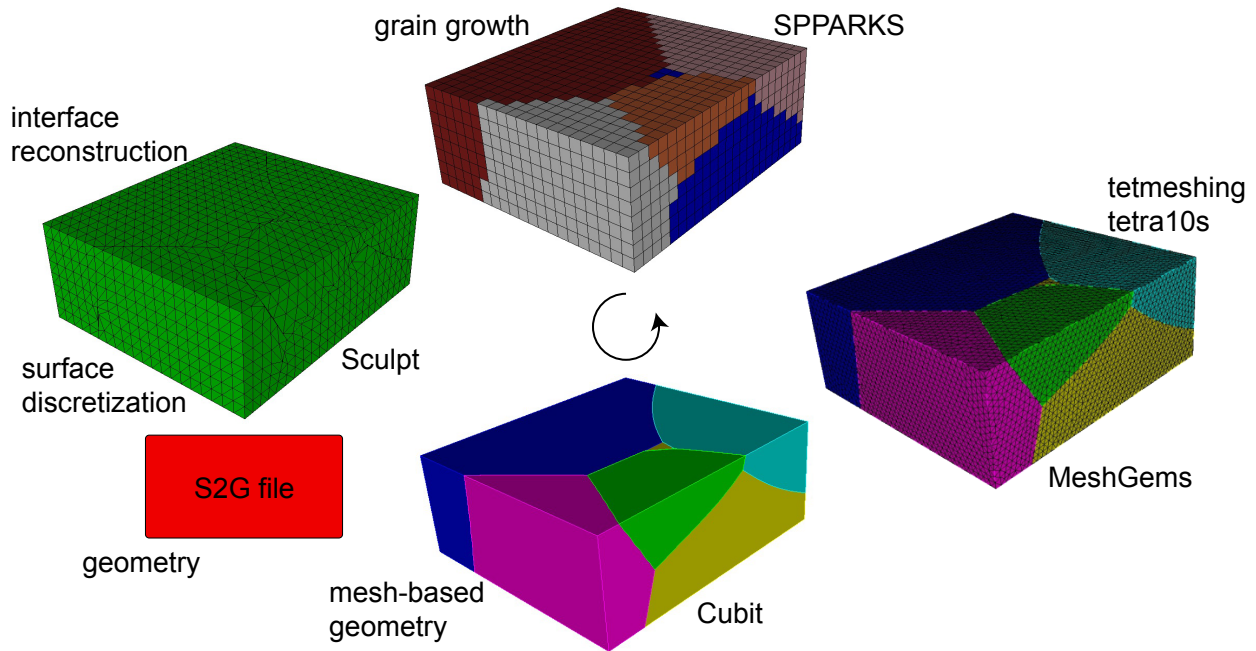


**Figure 24.** Polycrystalline microstructure generated from phase field grain growth simulations having 52 grains. (a) Phase field representation of the 3D grain microstructure. The scalar field  $\phi$  is used for visualization, where (red) blue denote (GBs) bulk grains. (b) FE discretization of the microstructure and (c) view of three grains in FEA mesh



**Figure 25.** Electron backscatter diffraction (EBSD) measurement and FE discretization using interface conformal hexahedral mesh.

vertices, curves, surfaces, and volumes. The workflow is illustrated in Figure 26. In an effort



**Figure 26.** Tetrahedral workflow for conformal boundaries that leverages Sculpt, Cubit, and MeshGems.

to illustrate the process, we consider the following process with components in grain growth, interface reconstruction, mesh-based geometry, and tetmeshing. To generate discretized, conformal microstructure we

1. Employ either SPPARKS or a phase field (PF) description to generate microstructure. If we employ the PF approach, we resolve a smooth partial differential equation on a particular grid and then voxelate the response.
2. Given that voxelated microstructure, employ Sculpt

```
sculpt -i spparks.i
```

with the input deck spparks.i having parameters

```
BEGIN SCULPT
  input_cart_exo = spparks.g
  trimesh = true
  exodus_file = spparks_tris
  write_geom = true
  defeature = 3
  smooth = no_surface_projections
  laplacian_iters = 3
END SCULPT
```



that write the *s2g* file, maximize defeaturing, limit the Laplacian smoothing to 3 iterations rather than the default 10. Although one may indeed desire more smoothing iterations, Laplacian smoothing will remove curvature. With shallow features near surfaces the possibility of creating slivers increases with additional smoothing. Two files result from interface reconstruction, *spparks\_tris.e.1.0* and *spparks\_tris.s2g*. Additional parameters in Sculpt can be invoked [79]. Those parameters and the defaults are exposed through the command line

```
sculpt --help
```

3. We now bring both the surface discretization, *spparks\_tris.e.1.0*, and the accompanying geometry file, *spparks\_tris.s2g*, into Cubit through

```
set developer on
import s2g "spparks_tris"
```

where we the import command looks for extensions *\*.e.1.0* and *\*.s2g* for the surface representation and geometry, respectively. We note that the current capability to import through an *s2g* representation is only available in serial.

4. Given a geometric definition of microstructure, we employ the MeshGems tetmesher in Cubit through

```
reset body all
volume all scheme tetmesh
trimesher surface gradation 1.2
trimesher volume gradation 1.2
set tetmesher optimize level 6 overconstrained off sliver on
set tetmesher interior points on
set tetmesher boundary recovery off
volume all size {mesh_size}
mesh volume all
```

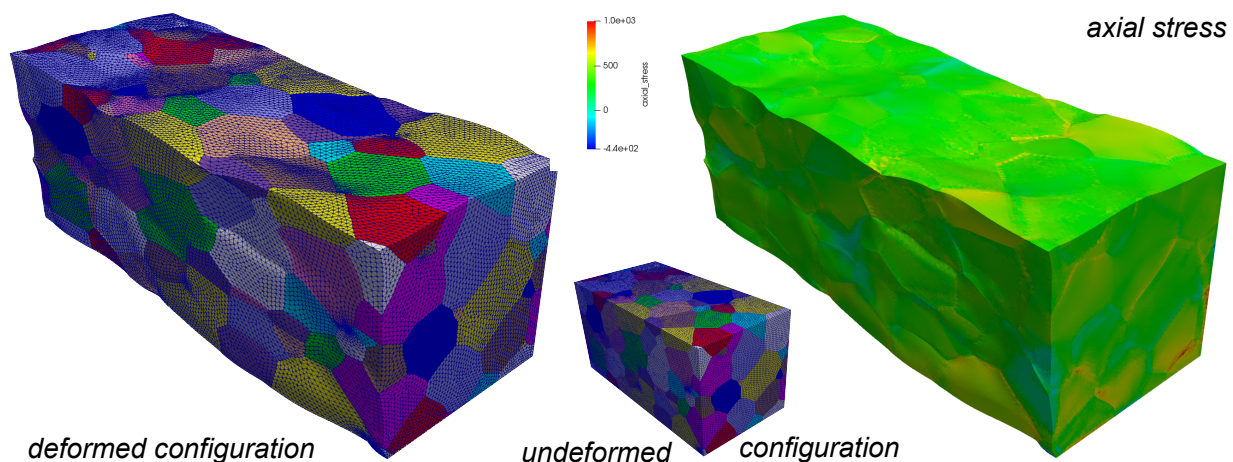
where the gradation factor is slightly less than the default, 1.3. In addition, the optimization level is maximized for element quality. Finally, one must declare the linear tetrahedral elements to be higher order tetrahedral elements through

```
block all element type tetra10
```

and then write an Exodus file for analysis.

In parallel with Figure 23, we can discretize candidate SPPARKS microstructures with higher-order tetrahedral elements amenable for the nearly isochoric motions associated with large plastic deformations. A particular configuration at  $\sim 0.45$  engineering strain is illustrated in Figure 27. We note that two planes of symmetry are enforced along the axis of loading which permit the loaded faces with normals parallel to the axial direction to contract. Although the applied boundary condition is one of plane stress, the anisotropy of the grains, as illustrated in prior sections, fosters a multi-axial response. These particular simulations employed slip-based hardening and were conducted in the Laboratory of Computational Mechanics (LCM). We achieved the illustrated configuration in 225 time steps. The simulation is not spellbinding. That said, we only seek to illustrate the utility of the recently formulated composite-tetrahedral element [56] for capturing


inhomogeneous crystal deformations and the boundary curvature that stems from random texture. We can continue to load the body to an engineering strain of nearly  $\sim 0.6$ . In particular grains, the



**Figure 27.** Conformal grain boundaries reconstructed from a grain-growth simulation and discretized with higher-order tetrahedral elements subjected to substantial deformations. A far-field engineering strain of 0.45 was achieved in 225 time steps. The volume-averaged axial stress and grain-boundary “ridges” highlight the grain-grain interactions.

deformation is localizing along grain boundaries. The resulting localization is noted in Figure 28. Clearly more work is needed to investigate local interactions and probe the evolution of localization with respect to both grain morphologies and grain texture. For this portion of the work, Corey Ernst focused on enabling a tetrahedral workflow for microstructure. To aid that process, Hojun Lim produced 10 microstructures akin to Figure 27 having on the order of 150 grains. The resulting microstructures ranged from 2 to 4 million tetrahedral elements. When Corey initiated this work in FY17, no microstructure would mesh in its entirety. By the end of the fiscal year, each of the 10 microstructures would not only mesh but was placed in a test repository to ensure that it would continue to “mesh.” We note that these developments were funded by a consortium of teams that focused on mesh-based geometry and with an emphasis on microstructure. Additional improvements to MeshGems were also accomplished through an NW LDRD devoted to the analysis of stronglinks. These improvements include

1. Integrated MeshGems MG-CADSurf into Cubit, for triangle meshing discrete surfaces (mesh-based/virtual geometry)
  - (a) Gradation control enabled on discrete geometry
  - (b) Triangle mesh size based on geometry approximation angle for discrete geometry
2. Resolved several MeshGems sizing problems



**Figure 28.** Illustrating the axial component of the total deformation gradient at and applied engineering strain of  $\sim 0.6$ . The deformation localizes along grain boundaries with local stretches are on the order of 9.



3. Fixed limitations on triangle meshing merged surfaces with MeshGems
4. Enhanced Sculpt to provide cleaner, smoother discrete geometry for successful tetmeshing

Through this work, we have enabled a general, tetrahedral workflow for microstructure. Our efforts complement prior work in hexahedral meshing and strengthen common capabilities in the generation and reconstruction of computational microstructure.

## 6 Structure-property linkages

We have established an Academic Alliance collaboration with Professor Kalidindi at Georgia Institute of Technology to understand microstructure-mechanical property relationships using a data science approach. In this work, we demonstrate the first application of the recently formulated Materials Knowledge Systems (MKS) framework in establishing reliable quantitative linkages between materials microstructures and their sensitivity to damage initiation [69]. Synthetic two-phase microstructures were generated as a calibration dataset with a wide range of microstructural features (including relative size of phases, morphology, spacing and eccentricity). The MKS framework was applied to obtain a reduced-order model for the structure-property linkage, which accounts for the volume averaged damage as well as the characteristics of the extremal distributions of the local damage in the heterogeneous microstructures.

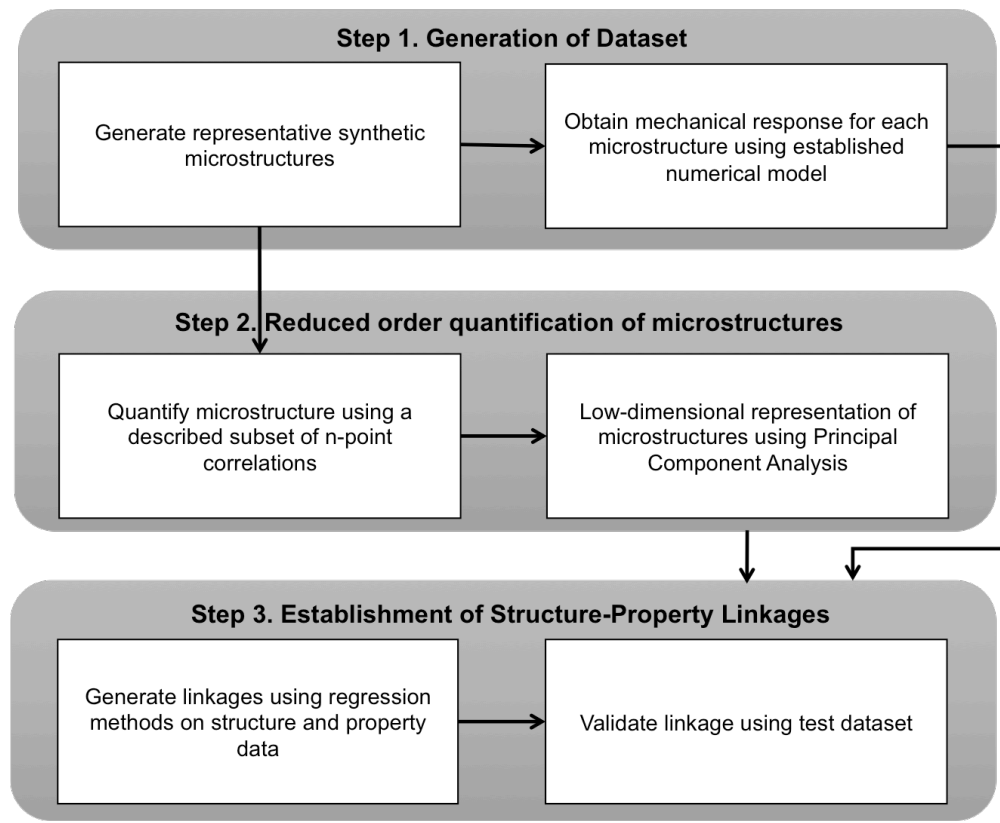
### 6.1 MKS framework

A typical workflow used in the implementation of the MKS framework for extracting structure-property linkages from physics-based simulation toolsets (e.g., micromechanical finite element models) is depicted in Figure 29 [27]. The approach starts with the generation of a dataset, where each data point constitutes one input (i.e., microstructure) and its corresponding output (i.e., the effective property of interest). The dataset is typically designed to include a large ensemble of microstructures of general interest to the specific phenomenon being studied. The microstructures are typically generated digitally and the mechanical response is evaluated numerically using established simulation tools (e.g., finite element models).

Material microstructures are quantified using the framework of  $n$ -point spatial correlations and obtaining data-driven (objective) low dimensional representations using principal component analysis (PCA). These PC representations of the microstructure (generally referred as PC scores) are then utilized as features in generating the reduced-order models connecting the microstructure to the property of interest. Various cross-validation techniques are employed to quantify the robustness and accuracies of the extracted models, and for selecting the best model. Details of model development including two-point statistics, PCA and cross-validation can be found elsewhere [Popova et al., 2017].

### 6.2 Dataset generation

The data generation for this study involves two main steps: (i) generation of a large ensemble of microstructures reflecting a rich diversity in microstructural features, and (ii) FE simulation of the damage initiation response for each microstructure in the ensemble using a suitable numerical simulation tool. A total of nine different classes of two-phase microstructures were generated using the different techniques described in Figure 30. Each class consisted of twenty different realizations; these were generated using nominally the same strategy but differ from other members of the same



**Figure 29.** Generalized MKS workflow for extracting structure-property linkage.

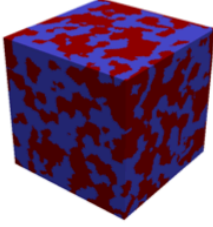
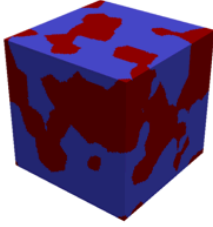
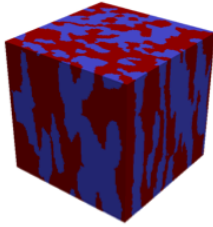
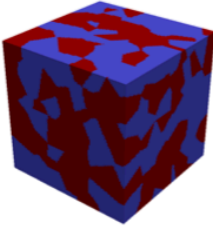
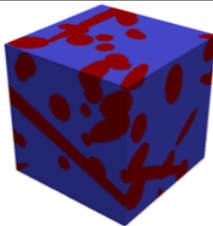
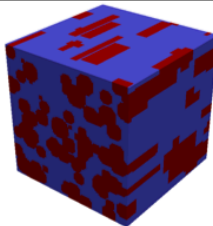
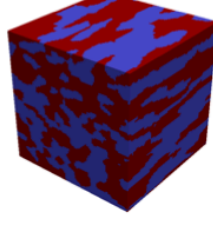
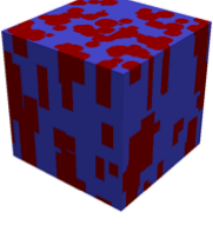
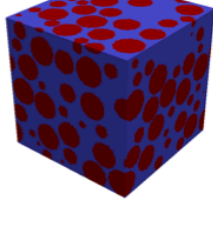

class in some stochastic manner. The volume fraction of each phase for all the microstructures was kept at  $\sim 50\%$ . As a result, the observed differences in the responses of the different microstructures included in this study are all attributed to the differences in the higher-order statistics describing the microstructure (note that the volume fraction information constitutes 1-point statistics of the microstructure). One example of a generated microstructure in each class is shown in Figure 30. The loading applied here is uniaxial tension along the Y-axis. It should be noted that Classes 7 and 8 are obtained by rotating microstructures of Classes 3 and 6, respectively, by 90 degrees. Therefore, in Class 6, the fibers are aligned with the loading direction (Y-axis), while in Class 8 the load is perpendicular to the orientation of the fibers. This allows exploration of the effect of anisotropic features on the damage response. The microstructures depicted in Figure 30 are discretized uniformly into  $100 \times 100 \times 100$  hexahedral elements.

In this work, finite element (FE) simulations of uniaxial tension were performed for each digitally generated microstructure, where an elastic-perfectly plastic (no hardening) material model is used for each constituent phase in the microstructure. Each phase was assigned a Young's modulus of 1000 GPa, a Poisson's ratio of 0.3, and yield strengths of 100 MPa (softer phase) and 500 MPa (harder phase), respectively. The softer phase is colored in red and the harder phase is shown in blue in Figure 30. The high value of Young's modulus was chosen to mitigate any effects from elasticity. Uniaxial tensile strain was applied along the Y-axis up to 1% overall (average) strain at a nominal strain rate of  $10^{-4} \text{ s}^{-1}$ . The contrast in the plastic properties produces heterogeneous stress and strain fields in the simulated microstructure that evolve significantly with the imposed overall strain of only 1%. Planar boundary conditions were applied to the microstructure, where one boundary plane was fixed in the Y-direction, and an appropriate displacement was prescribed on the opposite boundary plane (on the other side of the cube).

A common approach to simulating damage evolution in materials subjected to various loading conditions is to employ a void growth and coalescence model, where the initial porosity in the material is used to initialize the damage variable [28, 53, 82]. Such void nucleation and growth models can then be coupled with FE models to track damage evolution as a function of the applied loading conditions. These approaches typically employ suitable combinations of the equivalent plastic strain and stress triaxiality [10, 44, 46–49, 70], maximum principal tensile stress [18], hydrostatic stress [16], equivalent stress [17] as the main drivers of damage and its evolution. In this study, we have adopted a model based on plastic strain weighted cumulative triaxiality as a damage initiation indicator. Since this model accounts for both stress triaxiality and the imposed plastic strain, it captures some of the main established physics underlying damage initiation in ductile materials. This damage indicator [10, 49] is calculated as

$$D(\epsilon_{eq}) = \int_0^{\epsilon_{eq}} TF(\epsilon_{eq}) d\epsilon_{eq} \quad (61)$$

where  $TF$  is the stress triaxiality and  $\epsilon_{ep}$  is the equivalent plastic strain. In evaluating the sensitivity of the microstructures to damage initiation, the local damage parameter defined in Eq. (61) is evaluated for each voxel (element) in each microstructure. Although the treatment of the damage parameter is somewhat simplified in this work, our focus here is on the development of a general framework that allows consideration of a broad variety of definitions for the local damage

	Class 1. Polycrystalline microstructure with small equiaxed grains generated using DREAM.3D, and segmented into a 2-phase structure		Class 2. Polycrystalline microstructure with big equiaxed grains generated using DREAM.3D, and segmented into a 2-phase structure
	Class 3. Polycrystalline microstructure with elongated grains generated using DREAM.3D, and segmented into a 2-phase structure		Class 4 (test set). Two-phase microstructure generated using <u>Voronoi</u> tessellation
	Class 5. Randomly oriented fibers placed randomly in a matrix		Class 6. Aligned fibers placed randomly in a matrix
	Class 7. Class 3 rotated by 90 degrees along $x$ axis		Class 8. Class 6 rotated by 90 degrees along $x$ axis
	Class 9. Randomly placed spheres of randomly selected radii (within a specified range) in a matrix		

**Figure 30.** 2-phase microstructure classes used in this study along with the relevant details of how they were generated.

parameters in future studies.

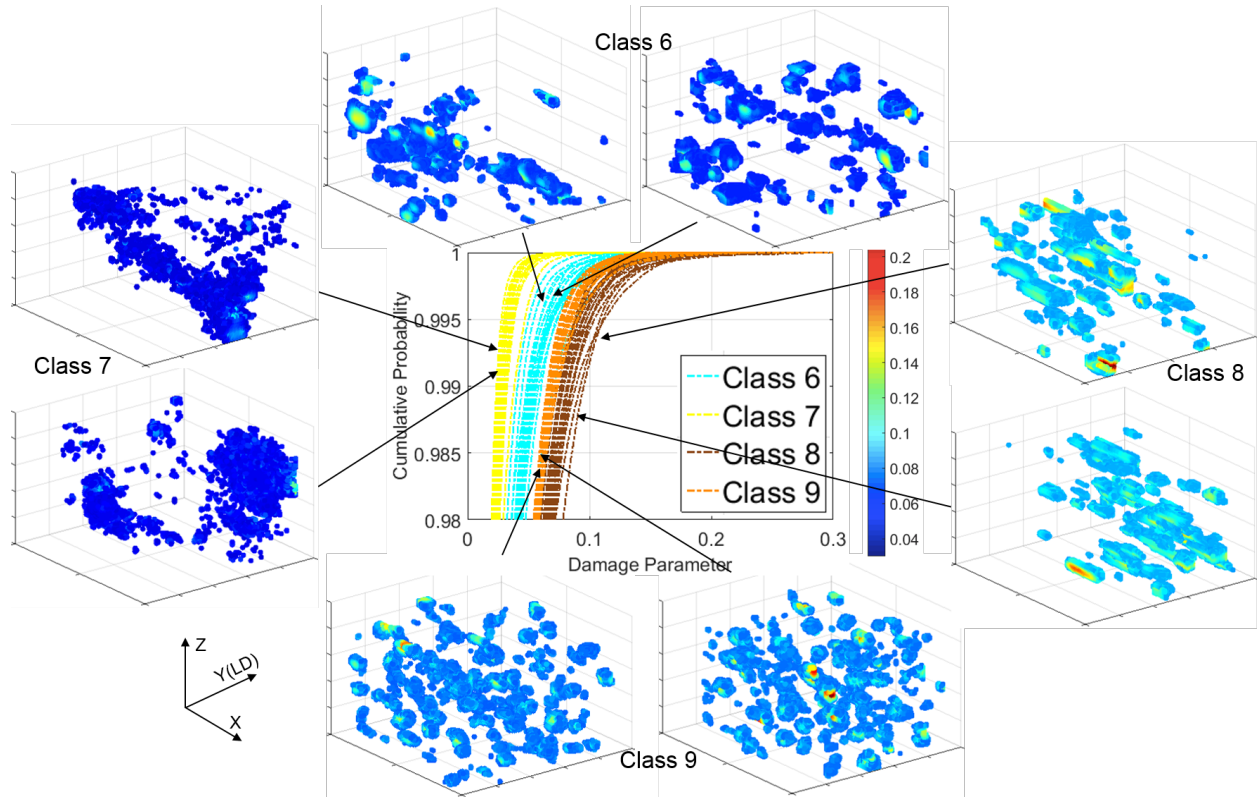
The cumulative distribution function (CDF) represents a probability that a random variable  $X$  will have a value equal or less than  $x$ , i.e.,  $F_X(x) = P(X \leq x)$ . The CDF of the local damage parameter values are obtained from the FE simulation for each microstructure. A Generalized Extreme Value Distribution (GEVD) function is then fit only to the top 2% of the CDF values to obtain the parameters  $k$ ,  $\sigma$  and  $\mu$  using MLE. This is done to define the effective damage parameter based on the extreme values of the local damage parameter. This is an important distinction for this effective damage measure, especially compared to the other simpler volume averaged damage measure used in this study. A GEVD function is a flexible three parameter distribution, whose Cumulative Distribution Function (CDF) is expressed as [34]

$$F(x; \mu, \sigma, k) = \exp \left\{ - \left[ 1 + k \left( \frac{x - \mu}{\sigma} \right) \right]^{-1/k} \right\}. \quad (62)$$

In Eq. (62),  $k$ ,  $\sigma$  and  $\mu$  denote the shape, scaling and location parameters of the GEVD function, respectively. Note that the top 2% produced robust values of the GEVD parameters that showed very little sensitivity to the choice of the CDF percentile cutoff value. Figure 31 demonstrates the top 2% of CDFs for all microstructures in Classes 6-9 (center of the Figure), and the contour plots depicting the spatial distribution of the top 2% of damage parameter for randomly chosen realizations of microstructures from Classes 6-9. Analysis of the CDF curves of damage parameter shows that the anisotropic microstructures that were loaded along the direction of the elongated phase (such as aligned cylinders (Class 6) and microstructures with elongated grains (Class 7)) demonstrate better performance (less sensitivity to damage initiation) than those loaded perpendicular to the interfaces (Class 8). It can be seen from the CDF plots in Figure 31 that microstructures corresponding to Class 7 perform the best, in terms of sensitivity to damage initiation, while microstructures of Class 8 perform the worst, i.e., more prone to damage initiation than other microstructures. When the CDFs are close to each other in Figure 31, the corresponding local damage parameter distributions are also similar, justifying the selection of this measure as an effective damage property of the microstructure. Furthermore, the contour plots depicted in Figure 3 show that the higher local damage parameter values are clustered along the interfaces on the side of the softer phase (Classes 6 and 8) or inside the softer matrix when the average spacing between the hard phase is small (e.g., spatially clustered spheres in microstructures of Class 9).

### 6.3 Reduced order model

The next step in the MKS framework is to establish a reduced-order structure-property linkage. The PC scores obtained in the previous section represent the microstructure. As described in the previous section, two different effective measures of damage were considered in this study: (i) the volume averaged damage, and (ii) the set of parameters describing the characteristics of the top 2% of the CDF for the local damage values. In the model building, the microstructure is represented through selected features. For example, the PC scores by themselves can be selected as features. In prior work, we have shown that polynomial products of the PC scores serve as excellent features in



**Figure 31.** CDF of damage parameter (top 2%) and contour plots of the top 2% of Damage parameter for Classes 6, 7, 8 and 9. Color scale bar is the same for all contour plots.

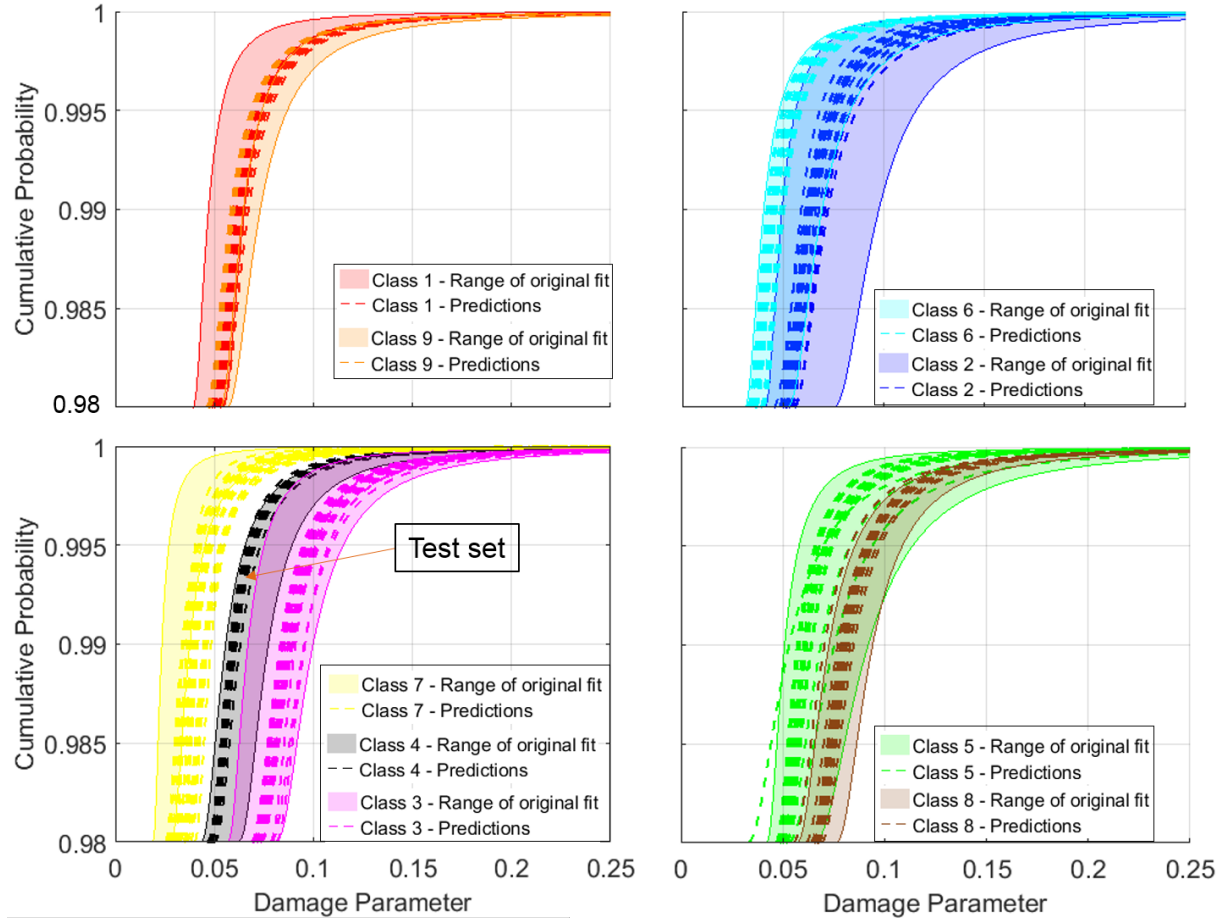
arriving at highly reliable PSP linkages [Latypov and Kalidindi, 2017]. In this work, polynomial products of the PC scores up to 4th order were considered in the model building process. The reduced order model was obtained using an ordinary least-squares regression technique.

In order to demonstrate the predictive capabilities of the reduced-order models developed in this work, the predicted GEVD parameters were used to reconstruct GEVD CDFs for all the microstructures. Figure 32 demonstrates the top 2% of CDF values for all the Classes of microstructures. As expected, there are significant overlaps in the predicted CDFs (these were also present in the original dataset extracted directly from the FE models). For convenience of viewing, the plots are presented as four groups in Figure 32. Each shaded area in each plot represents a range of the original GEVD fit for each microstructure Class, while the dashed lines are GEVDs from the predicted values of  $\mu$  and  $\sigma$  using the reduced-order models extracted in this work. Predictions for most Classes fall into the range of the original fits. Overall, the errors are acceptable for the current application of rank-ordering sensitivity to damage initiation.

## 6.4 Summary

In this work, the MKS framework was extended and applied successfully for the first time to rank-ordering sensitivity of two-phase composite microstructures to damage initiation. The framework was used to capture the complex, and not readily discernible connections between the many features of a microstructure and its sensitivity to damage initiation. Two measures of damage initiation were considered in the current work. The first approach considered volume averaged values of a local damage parameter as an effective property; the second approach used GEVD parameters that were fit to the distribution of the top 2% of the values of the damage parameter. While the first approach offers a simple measure, the second aims to capture some aspects of the extreme-value distributions associated with the damage initiation process. All models exhibited a remarkable ability to capture the dominant aspects of the very complex structure-property linkages in the microstructures studied. As expected, the model for the average values of damage parameter produced less error compared to the models of GEVD parameters. However, the models for GEVD parameters provide more insight into the distribution of the local damage parameter. The successful application of the MKS framework to the highly complex problem studied here demonstrates its versatility. It is indeed remarkable that the MKS framework is able to identify objectively from the provided data the main microstructural features and their quantitative influence on the damage initiation parameters. This is of particular importance for future studies aimed at establishing PSP linkages in heterogeneous materials, as this approach does not require any ad-hoc selection of specific microstructure measures.





**Figure 32.** Top 2% of CDF of the original fit to GEVD (shaded area) and GEVD from predicted parameters (dashed lines) for all of the microstructures.

## 7 Discussion and future work

The work presented herein represents attempts to outline a framework for investigating mechanisms through the simulation of microstructure. We review our findings, reflect on lessons learned, and plot a path for future development in each effort.

### 7.1 Requisite components

#### Concurrent multiscale

In this study, a multiscale simulation of a tensile specimen was performed. The microstructure of the specimen material was represented by an explicit discretization of the polycrystalline grain structure in the gauge section of the specimen. Outside of this region, a much coarser discretization was employed to capture only the geometry of the specimen. Local heterogeneity in the material was not reflected in the simulation. The microstructural subdomain was linked to two specimen-scale subdomains via the Schwarz alternating method of concurrent multiscale coupling. It was demonstrated that this method could be used to investigate the role of microstructure on the heterogeneous deformation of the specimen in a computationally tractable manner.

Beyond this demonstration, the simulation framework developed in this study will be used to analyze localization during necking of the tensile specimen and the impact of the microstructure on this process. A number of key developments have been identified during the present study that are required to enable this exploration. First, the discretization of the microstructural domain must be evolved to conform to the grain boundaries in the microstructure if stress concentrations are to be modeled properly. Second, the role of the material model in the specimen-scale subdomains should be examined in order to determine a more commensurate representation of the aggregate microstructural response with respect to anisotropy and rate dependence. Lastly, the impact of the Schwarz coupling of different subdomains with disparate material models must be evaluated.

#### Modular approach to crystal plasticity

1. *Additional physics.* Although we did make progress in many aspects of the flow of single crystals, we did not make fundamental contributions in deformation twinning. We believe that twinning is an important mechanism in the localization of the deformation for hydrogen-charged stainless steel alloys and tantalum alloys exposed to elevated rates of loading. To better understand the importance of twinning on the deformation of stainless steel alloys, we have proposed a set of experiments probing single crystals noted in Section 7.2.
2. *Additional numerics.* These remarks pertain to the numerical infrastructure that we developed for the crystal plasticity model. For discussion about the numerical techniques employed for the solution of the global systems of equations, refer to the Schwarz alternating

method used for concurrent multiscale. The use of MiniSolver and MiniTensor for the implementation of the crystal plasticity model provided many features, such as support for manycore architectures, and led to significant reduction in the complexity of the implementation. Also, early in the project there was evidence that the use of MiniSolver provided a decrease in computation time of up to 30%. The crystal plasticity model is significantly complex, and challenges any numerical solution scheme. Although MiniSolver became the work horse for crystal plasticity calculations, it is clear that some new features would improve both its robustness and efficiency. Among these are: improving the numerical efficiency of constrained optimization algorithms, use of sub-cycling, and adaptive use of solution schemes (e.g. starting with one scheme and switching to another later on). The crystal plasticity model was the motivation of many features introduced to MiniSolver, and it tested the limits of its design.

3. *Crystal plasticity in Sierra/SolidMechanics.* As described in Section 4.4, the crystal plasticity constitutive model has been ported to LAME for use in Sierra/SolidMechanics. This is a crucial element in advancing the technology developed in this project from the research and development phase to the production phase, where it can be used directly for the fulfillment of Sandia's nuclear weapons mission.

Successful deployment of the crystal plasticity model to the Sandia analyst community relies on 1) the ability to support analysts in the application of the material model to mission-relevant engineering simulations, and 2) the ability to perform ongoing software maintenance on the material model source code. It is clear, based on preliminary testing, that successful use of the crystal plasticity model in complex engineering simulations will require close collaboration between the model developers and analysts using the Sierra code suite. Collaboration is critical because exercising the material model in this context reveals shortcomings in the implementation which must be addressed by a model developer, for example in the treatment of numerical difficulties that can arise under conditions of large deformation and high strain rates. Continued involvement of the model development team is also necessary in the context of standard, ongoing software maintenance. Sandia engineering codes are subject to a process of continual improvement that can necessitate changes to material model code. This is true within the Sierra codes themselves, and also within the LCM code, where changes to the crystal plasticity model may need to be ported to the LAME implementation.

Several activities planned for FY18 are congruent with the present LDRD and provide a means for advancing the crystal plasticity model and expanding its impact. The multiscale technology demonstrator being developed under ASC/ATDM serves as an early adopter for next-generation software tools, e.g., the DARMA asynchronous many-task scheduler and the Kokkos library for performance portability. A FE<sup>2</sup> multiscale capability was added to the technology demonstrator in FY17 and provides a natural application space for the crystal plasticity model. In this context, the crystal plasticity model can be used within a mesoscale representative volume element, capturing the effects of grain-scale mechanisms and informing the overall macroscale response. A second FY18 activity in which the crystal plasticity model is highly relevant is the ASC/P&EM project on adapting LAME material models for emerging, next-generation computing hardware. Significant performance gains for the crystal plasticity model are possible under this project, which could greatly improve usability in

engineering simulations.


4. *Manycore architectures.* There was significant progress in the introduction of manycore parallelism for the crystal plasticity model by means of the Kokkos framework. This is not a trivial task, and despite technical and resource availability challenges, we succeeded. But there is still much work to do. Specifically, we only tested OpenMP parallelism. Kokkos supports other types of devices that implement different modes of parallelism, such as CUDA or Intel Phi. We did not have access to the necessary infrastructure to test our implementation in these platforms, however. In addition, the finite element and solver infrastructure needs to be parallelized to realize full speedups. That task, however, is well beyond the scope of this work.
5. *Robustness suite.* Although progress was made in the conception and execution of a robustness suite, we need cluster builds to be both diverse and stable to support a large suite of performance tests. The Robustness suite should have sufficient breadth to provide a technology pull for better methods. In addition to providing feedback on the applicability of new algorithms to a suite of increasing complexity, we also seek to provide automation in the generation, submission, simulation, and reduction of ensembles of microstructures.

## Conformal microstructures

We have made significant progress in moving towards conformal approaches and away from voxelated representations during the last two years of this project. We have now able to examine both a hexahedral workflow and tetrahedral workflow for microstructure. Throughout this LDRD, we have deployed resources to enable the tetrahedral workflow because we believe there are inherent advantages to tetrahedral workflow. We can support 3D size fields with tetrahedral elements. We also have a surface representation that is independent of the chosen discretization which permits mesh convergence studies to assess both near-field and far-field information.


There are, however, a few issues that merit increased scrutiny and new methods. With regard to the phase field approach, we currently solve a 4<sup>th</sup>-order partial differential equation, voxelate the continuous solution, and then reconstruct the voxelated approximation. Given the limitation of voxelated input, we might propose to mitigate error through recursive bisection. Figure 33 attempts to illustrate this simple concept in 2D. More recent work towards the end of FY17 in Sculpt has also investigated a space-filling tetrahedral overlay rather than a hexahedral overlay. While this may indeed provide additional opportunities for mitigating error, we need targeted efforts to accurately map PF simulations for geometric reconstruction. We need to take advantage of the smooth isosurfaces (0.5) inherent in the phase field approach.

Another issue inherent in grain-growth simulations stems from boundary conditions. Because the domain sizes are relatively small, periodic boundary conditions are typically employed for grain growth. Albeit logical, there are ramifications for analysis efforts. Triple junctions and high aspect ratio structures may be located adjacent to one boundary while the majority of the structure is connected to the opposite surface. The small features that result from this process will drive down the element size needed for resolution. Figure 34 attempts to illustrate these matters for a



**Figure 33.** Schematic showing different levels of adaptive mesh refinement via recursive bisection. The resulting octree data structure provides increased accuracy to resolve grain boundary curvature and triple junctions.

particular microstructure. Please note that the highly discretized regions are driven by these fine features. Some of the finest feature are already being altered by defeaturing. More study is needed. As noted in boundary conditions applied in Sections 2 and 3, periodic boundary conditions may not always be appropriate. We should consider additional boundary conditions in our approach to grain growth. From a practical standpoint, we will always section our experimental observations of microstructure. Given a sufficient number of grains, our representations will contain these features. Thoughtful study is required to focus on mechanisms and the relevance to these particular features.



**Figure 34.** Cases of substantial refinement that may stem from periodic boundary conditions in the grain growth simulations. More study is needed to understand the relevance of these features.

## Structure-property linkages

An academic alliance with Surya Kalidindi's group at Georgia Institute of Technology enabled the extension of MKS framework to investigate two-phase microstructures to damage evolution. This framework was used to capture the relationship between the many features of a microstructure and its sensitivity to damage initiation. The next step would be to incorporate more detailed microstructural features, such as crystal orientations of polycrystals and their deformation behaviors. This would require a large number of simulations for training data. Furthermore, various damage initiation and evolution criteria should be tested and validated with experiments.

In addition, we will continue to collaborate with Surya Kalidindi to develop a database-driven crystal plasticity framework. This method employs a spectral (Fourier) representation of the responses of an efficiently selected number of crystal plasticity simulations to develop a database that allows for the rapid calculation of the crystal plasticity constitutive equations. This new technique calculates the results of crystal plasticity simulations 100 times faster.

## 7.2 A path forward through teaming

The work initiated in this project will move forward next year on many fronts. We will pursue aspects of this problem in Laboratory Directed Research & Development (LDRD), Physics & Engineering Models (P&EM), Advanced Technology Deployment & Mitigation (ATDM), Campaign 2 (C2), Modeling Materials Assurance (MMA), Weapons Systems Engineering Assessment Technology (WSEAT), and University partnerships. FY18 projects include

1. Uncertainty quantification of microstructural material variability (Jones, Alleman, ES LDRD)
2. Dynamic strain aging in additive manufactured alloys (Antoun, Alleman, NW LDRD)
3. Upscaling mesoscale mechanisms/uncertainty (Foulk, Alleman, Emery, P&EM)
4. Schwarz multiscale coupling (Mota, Emery, P&EM)
5. Microstructural mechanisms of H embrittlement (Foulk, Alleman, P&EM)
6. Constitutive models for next generation architectures (Littlewood, P&EM)
7. Using DARMA for FE<sup>2</sup> asynchronous demonstrator (Littlewood, ATDM)
8. Materials Margin Assurance (MMA), ARMS springs (Madison, Lim, P&EM)
9. High rate deformation of tantalum (Lane, Lim, C2)
10. Modeling two-phase Ni-based superalloys (Lim, Cambridge)
11. Continue to probe structure-property relations for localization (Lim, GT)



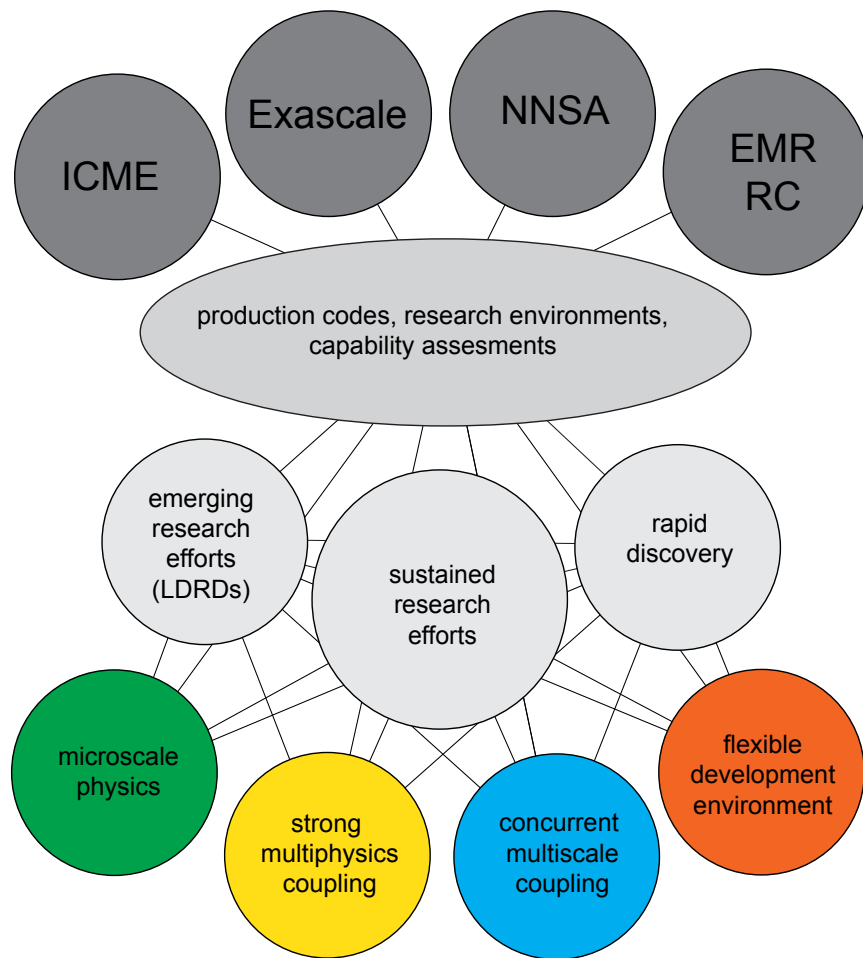
Given that we are now employing tools to solve classes of problems, we must be very deliberate in teaming to develop fundamental capabilities for the community. Moreover, additional efforts are needed to support both the aforementioned projects and our mission objectives. We propose to fund following projects

1. Influence of hydrogen isotopes on mechanical and microstructural evolution of discrete microstructural elements at dislocation length scales (San Marchi, WSEAT)
2. Mesoscale phase field models of microstructure formation and evolution using SPPARKS (Abdeljawad)

to provide fundamental experiments of single crystals and new methods for generating relevant computational microstructures. Mesoscale modeling is built on the flow of single crystals. We cannot proclaim modeling and simulation to be relevant for our understanding without conducting experiments to validate the “building blocks” of our approach. In that vein, we cannot assert that microstructure is the key to understanding mechanisms if we are not going to invest in computational microstructure. None of the 11 aforementioned projects in FY18 has a major component in generating relevant microstructures for computation. Again, we must be deliberate in our approaches and share those findings with the community.

It is our hope that we can continue to contribute in the areas of microscale physics, strong multiphysics coupling, concurrent multiscale coupling, and a flexible development environment. As shown in Figure 35, we believe that these fundamental contributions enable sustained research efforts, spawn emerging efforts, and provide an opportunity for rapid discovery. The totality of these efforts implemented in research and production environments will support of internal initiatives (Engineering of Materials Reliability Research Challenge), enable stockpile stewardship (National Nuclear Security Administration), and promote collaboration through national (Exascale) and international (Integrated Computational Materials Engineering) initiatives.





**Figure 35.** Components of a technical basis that enable research efforts, mature research environments, support production implementation, and ultimately serve our mission and foster collaboration with national/international communities.

## 8 Conclusion

Through this work, we attempted to push forward in key areas to connect microscale physics to structural performance.

1. We were successful in simulating the two-way coupling of a heterogeneous polycrystal employing crystal plasticity to a homogenous bar employing  $J_2$  plasticity using the Schwarz Alternating method. Studies focused on incipient localization and the role of boundary conditions applied to microstructure.
  - (a) Performed homogenization for the  $J_2$  overlap region
  - (b) Contrasted multiple variations of the Schwarz Alternating method in LCM
  - (c) Recovered monotonic convergence with the classic Schwarz Alternating method
  - (d) Focused on verification, test-driven development, and robustness
2. We have developed a modular approach to crystal plasticity in which one can employ the desired physics and the requisite numerics for implicit solution.
  - (a) Implemented modular physics for the flow rule and hardening laws
  - (b) Enabled modular numerics through the development of MiniSolver and MiniTensor
  - (c) Focused on verification, test-driven development, and robustness
3. Detailed studies quantified the complexity of the crystal physics. Novel diagnostics were developed to illustrate slip activity on all 12 slip systems and the impact on the residual and convergence rate during the iterative process.
4. Verification problems were developed to idealize fully-populated slip systems. The quad-slip case provides an opportunity for testing major components of the algorithmic update.
5. We accept that developing in research environment is not sufficient for widespread adoption. In an effort to support the production environment, we ported the crystal plasticity model to *Sierra/SM* multiple times during our development process.
6. We successfully illustrated the scalability of the crystal plasticity model for new architectures. Specifically, we achieved a speedup of  $\sim 20$  with 32 cores using the Kokkos OpenMP parallel mode.
7. In reflection of our values, we developed a robustness suite to exercise our algorithms on problems of increasing complexity. We believe that such a suite is necessary to quantify and monitor the performance of our algorithms.
8. Building on our success with generating conformal grain boundaries for a hexahedral workflow, we extended the generation of conformal boundaries to a tetrahedral workflow.
  - (a) Resolving the grain-grain interactions that facilitate strain localization

- (b) Demonstrated that composite tetrahedral elements provide resolution and smoothness
  - (c) Simulated large deformations (stretches  $\sim 9$ ) with implicit time integration
9. In collaboration with Professor Kalidindi at Georgia Tech, we applied Materials Knowledge Systems to synthetic two-phase microstructures. We linked Principal Component Analysis (PCA) of the microstructures with Generalized Extreme Value Distributions (GEVD) of void evolution through multivariate polynomial regression to link microstructural features to material performance.
  10. Lessons learned and future work were illustrated for the requisite components which include: concurrent multiscale, crystal plasticity, conformal microstructures, and structure-property linkages. A credible path forward for fundamental work requires teaming.

## References

- [1] F. Abdeljawad and S. M. Foiles. Stabilization of nanocrystalline alloys via grain boundary segregation: A diffuse interface model. *Acta Mater.*, 101:159–171, 2015.
- [2] F. Abdeljawad, B. Voelker, R. Davis, R. M. McMeeking, and M. Haataja. Connecting microstructural coarsening processes to electrochemical performance in solid oxide fuel cells: An integrated modeling approach. *J. Power Sources*, 250:319–331, 2014.
- [3] D. Abraham and C. Alstetter. Hydrogen-enhanced localization of plasticity in an austenitic stainless steel. *Metallurgical and Materials Transactions A*, 26:2859–2871, 1995.
- [4] C. N. Alleman, D. J. Luscher, C. A. Bronkhorst, and S. Ghosh. Distribution-enhanced homogenization framework and model for heterogeneous elasto-plastic problems. *J. Mech. Phys. Solids*, 85:176–202, 2015.
- [5] C. N. Alleman, J. W. Foulk, A. Mota, H. Lim, and D. J. Littlewood. Concurrent multiscale modeling of microstructural effects on localization behavior in finite deformation solid mechanics. *Computational Mechanics*, page accepted, 2017.
- [6] S. M. Allen and J. W. Cahn. A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening. *Acta Metall.*, 27:1085–1095, 1979.
- [7] F. Armero and A. Pérez-Foguet. On the formulation of closest-point projection algorithms in elastoplasticity—part i: The variational structure. *International Journal for Numerical Methods in Engineering*, 53:297–329, 2002.
- [8] F. Armero and J. Simo. A new unconditionally stable fractional step method for non-linear coupled thermomechanical problems. *International Journal for numerical methods in Engineering*, 35:737–766, 1992.
- [9] F. Armero and J. Simo. A priori stability estimates and unconditionally stable product formula algorithms for nonlinear coupled thermoplasticity. *International Journal of Plasticity*, 9:749–782, 1993.
- [10] Y. Bao and T. Wierzbicki. On fracture locus in the equivalent strain and stress triaxiality space. *International Journal of Mechanical Sciences*, 46(1):81–98, 2004.
- [11] N. R. Barton, N. W. Winter, and J. E. Reaugh. Defect evolution and pore collapse in crystalline energetic materials. *Modelling and Simulation in Materials Science and Engineering*, 17(3):035003, 2009.
- [12] Z. Bažant. Can multiscale-multiphysics methods predict softening damage and structural failure? *International Journal for Multiscale Computational Engineering*, 8:61–67, 2010.
- [13] F. B. Belgacem. The mortar finite element method with lagrange multipliers. *Numerische Mathematik*, 84(2):173–197, 1999.

- [14] T. Belytschko, S. Loehnert, and J.-H. Song. Multiscale aggregating discontinuities: A method for circumventing loss of material stability. *International Journal for Numerical Methods in Engineering*, 73:869–894, 2008.
- [15] M. Borrel, L. Halpern, and J. Ryan. Euler/Navier-Stokes couplings for multiscale aeroacoustic problems. *20th AIAA Computational Fluid Dynamics Conference, Honolulu, Hawaii*, 2011–3047, 6 2011.
- [16] P. Brozzo, B. Deluca, and R. Rendina. A new method for the prediction of formability limits in metal sheets. In *Proc. 7th biennial Conf. IDDR*, 1972.
- [17] S. E. Clift, P. Hartley, C. E. N. Sturgess, and R. G. W. Fracture prediction in plastic deformation processes. *International Journal of Mechanical Sciences*, 32:1–17, 1990.
- [18] M. Cockcroft and D. Latham. Ductility and the workability of metals. *J Inst Metals*, 96:33–39, 1968.
- [19] E. W. C. Coenen, V. G. Kouznetsova, and M. G. D. Geers. Novel boundary conditions for strain localization analyses in microstructural volume elements. *Int. J. Numer. Methods Eng.*, 90(1):1–21, 2012.
- [20] R. de Borst. Damage, material instabilities, and failure. In E. Stein, R. de Borst, and T. Hughes, editors, *Encyclopedia of Computational Mechanics*. John Wiley & Sons, Ltd., New York, 2004.
- [21] I. Demeshko, W. Spitz, I. Tezaur, J. Watkins, O. Guba, A. Salinger, R. Pawlowski, and M. Heroux. Towards performance-portability of the albania finite element analysis code using the kokkos library. *International Journal of High Performance Computing Applications*, page submitted, 2016.
- [22] C. Di Leo and L. Anand. Hydrogen in metals: A coupled theory for species diffusion and large elastic-plastic deformations. *International Journal of Plasticity*, 43:42–69, 2013.
- [23] R. Dingreville, C. Battaile, L. Brewer, E. Holm, and B. Boyce. The effect of microstructural representation on simulations of microplastic ratcheting. *International Journal of Plasticity*, 26:617–633, 2010.
- [24] M. Engel and M. Griebel. Flow simulation on moving boundary-fitted grids and application to fluid-structure interaction problems. *International Journal for Numerical Methods in Fluids*, 50(4):437–468, 2005.
- [25] C. Farhat, M. Lesoinne, P. LeTallec, K. Pierson, and D. Rixen. Feti-dp: a dual-primal unified feti method-part i: A faster alternative to the two-level feti method. *International Journal for Numerical Methods in Engineering*, 50(7):1523–1544, 2001.
- [26] F. Feyel and J.-L. Chaboche.  $FE^2$  multiscale approach for modelling the elastoviscoplastic behaviour of long fibre SiC/Ti composite materials. *Comput. Methods Appl. Mech. Eng.*, 183(3?4):309 – 330, 2000.

- [27] A. Gupta, A. Cecen, S. Goyal, A. K. Singh, and S. R. Kalidindi. Structure-property linkages using a data science approach: Application to a non-metallic inclusion/steel composite system. *Acta Materialia*, 91:239–254, 2015.
- [28] A. Gurson. Continuum theory of ductile rupture by void nucleation and growth: PART I. *Journal of Engineering Materials and Technology*, 99:2–15, 1977.
- [29] N. Hadjiconstantinou and A. Patera. Heterogeneous atomistic-continuum representations for dense fluid systems. *International Journal of Modern Physics C*, 8(4):967–976, 1997.
- [30] M. Heroux. An overview of trilinos. Technical Report SAND2003-2937, Sandia National Laboratory, 2003.
- [31] R. Hill. Elastic properties of reinforced solids: Some theoretical principles. *J. Mech. Phys. Solids*, 11(5):357–372, 1963.
- [32] R. Hill and J. R. Rice. Constitutive analysis of elastic plastic crystals at arbitrary strain. *J. Mech. Phys. Solids*, 20:401–413, 1972.
- [33] E. A. Holm and C. C. Battaile. The computer simulation of microstructural evolution. *JOM*, 53(9):20–23, 2001.
- [34] J. Hosking. Algorithm as 215: Maximum-likelihood estimation of the parameters of the generalized extreme-value distribution. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 34:301–310, 1985.
- [35] J. Hutchinson. Bounds and self-consistent estimates for creep of polycrystalline materials. *Proceedings of the Royal Society of London*, 348:101–127, 1976.
- [36] J. W. Hutchinson. Bounds and self-consistent estimates for creep of polycrystalline materials. *Proc. R. Soc. Lond. A*, 348:101–127, 1976.
- [37] K. Karlson, J. Foulk III, A. Brown, and M. Veilleux. Sandia fracture challenge 2: Sandia california’s modeling approach. *International Journal of Fracture*, 198:179–195, 2016.
- [38] S. Keshavarz and S. Ghosh. Hierarchical crystal plasticity fe model for nickel-based superalloys: Sub-grain microstructures to polycrystalline aggregates. *Int. J. Solids Struct.*, 55:17–31, 2015.
- [39] V. Kouznetsova, W. A. M. Brekelmans, and F. P. T. Baaijens. An approach to micro-macro modeling of heterogeneous materials. *Comput. Mech.*, 27(1):37–48, 2001.
- [40] E. H. Lee. Elastic-plastic deformation at finite strains. *Appl. Mech.*, 36:1–6, 1969.
- [41] H. Lim, C. R. Weinberger, C. C. Battaile, and T. E. Buchheit. Application of generalized non-Schmid yield law to low temperature plasticity in bcc transition metals. *Model. Simul. Mater. Sci. Eng.*, 21:045015, 2013.

- [42] H. Lim, J. Carroll, C. Battaile, T. Buchheit, B. Boyce, and C. Weinberger. Grain-scale experimental validation of crystal plasticity finite element simulations of tantalum oligocrystals. *International Journal of Plasticity*, 60:1–18, 2014.
- [43] H. Lim, F. Abdeljawad, S. J. Owen, B. Hanks, J. Foulk, and C. Battaile. Incorporating physically-based microstructures in materials modeling: Bridging phase field and crystal plasticity frameworks. *Modelling and Simulation in Materials Science and Engineering*, 24:045016, 2016.
- [44] A. C. Mackenzie, J. W. Hancock, and D. K. Brown. On the influence of state of stress on ductile failure initiation in high strength steels. *Engineering fracture mechanics*, 9:167–168, IN13–IN14, 169–188, 1977.
- [45] K. Matous and A. Maniatty. Multiscale modeling of elasto-viscoplastic polycrystals subjected to finite deformations. *International Journal for Numerical Methods in Engineering*, 60: 2313–2333, 2004.
- [46] F. McClintock. A criterion for ductile fracture by growth of holes. *Journal of Applied Mechanics*, 35:363–371, 1968.
- [47] F. McClintock. Local criteria for ductile fracture. *International Journal of Fracture Mechanics*, 4:101–130, 1968.
- [48] G. Mirone. Role of stress triaxiality in elastoplastic characterization and ductile failure prediction. *Engineering Fracture Mechanics*, 74(8):1203–1221, 2007.
- [49] G. Mirone. Elastoplastic characterization and damage predictions under evolving local triaxiality: axisymmetric and thick plate specimens. *Mechanics of Materials*, 40:685–694, 2008.
- [50] P. R. Morris. Elastic constants of polycrystals. *Int. J. Eng. Sci.*, 8(2):49–61, 1970.
- [51] A. Mota, I. Tezaur, and C. Alleman. The schwarz alternating method in solid mechanics. *Computer Methods in Applied Mechanics and Engineering*, page in press, 2017.
- [52] C. Mullner, C. Solenthaler, P. Uggowitzer, and M. Speidel. Brittle fracture in austenitic steel. *Acta Metallurgica et Materialia*, 42:2211–2217, 1994.
- [53] A. Needleman and V. Tvergaard. An Analysis of Ductile Rupture in Notched Bars. *Journal of the Mechanics and Physics of Solids*. *J. Mech. Phys. Solids*, 32(6):461–490, 1984.
- [54] K. Nibur, B. Somerday, D. Balch, and C. San Marchi. The role of localized deformation in hydrogen-assisted crack propagation in 21Cr-6Ni-9Mn stainless steel. *Acta Materialia*, 57: 3795–3809, 2009.
- [55] J. Nocedal and S. Wright. *Numerical Optimization*. Springer-Verlag, New York, 2006.
- [56] J. Ostien, J. Foulk, A. Mota, and M. Veilleux. A 10-node composite tetrahedral finite element for solid mechanics. *International Journal for Numerical Methods in Engineering*, 107(13): 1145–1170, 2016.

- [57] S. Owen, M. Staten, and M. Sorensen. Parallel hexahedral meshing from volume fractions. *Engineering with Computers*, 30:301, 2014.
- [58] S. J. Owen. Parallel smoothing for grid-based methods. In *21st International Meshing Roundtable, Research Note*, 2012.
- [59] S. J. Owen and T. R. Shelton. Validation of grid-based hex meshes with computational solid mechanics. In *Proc. 22nd International Meshing Roundtable*, pages 39–56, 2013.
- [60] S. J. Owen and R. M. Shih. A template-based approach for parallel hexahedral two-refinement. In *Proc. 24th International Meshing Roundtable*, pages 31–43, 2015.
- [61] S. J. Owen, M. L. Staten, and M. C. Sorensen. Parallel hex meshing from volume fractions. In *Proc. 20th International Meshing Roundtable*, pages 161–178, 2011.
- [62] S. J. Owen, J. A. Brown, C. Ernst, H. Lim, and K. Long. Hexahedral mesh generation for 3d computational materials modeling. In *Proc. 26th International Meshing Roundtable*, 2017.
- [63] V. Pandurangan, H. Li, and T. Ng. A concurrent multiscale method based on the alternating Schwarz scheme for coupling atomic and continuum scaled with first-order compatibility. *Computational Mechanics*, 47:1–16, 2011.
- [64] D. Paquet, P. Dondeti, and S. Ghosh. Dual-stage nested homogenization for rate-dependent anisotropic elasto-plasticity model of dendritic cast aluminum alloys. *Int. J. Plast.*, 27(10): 1677–1701, Oct. 2011. ISSN 07496419.
- [65] M. Parks, P. Bochev, and R. Lehoucq. Connecting atomistic-to-continuum coupling and domain decomposition. *Multiscale Modeling & Simulation*, 7(1):362–380, 2008.
- [66] R. Pawlowski, E. Phipps, and A. Salinger. Automating embedded analysis capabilities and managing software complexity in multiphysics simulation, part i: Template-based generic programming. *Scientific Programming*, 20:197–219, 2012.
- [67] R. Pawlowski, E. Phipps, and A. Salinger. Automating embedded analysis capabilities and managing software complexity in multiphysics simulation, part ii: Application to partial differential equations. *Scientific Programming*, 20:327–345, 2012.
- [68] D. Peirce, R. J. Asaro, and A. Needleman. An analysis of nonuniform and localized deformation in ductile single crystals. *Acta Metall.*, 30:1087–1119, 1982.
- [69] E. Popova, D. Montes de Oca Zapiain, F. Abdeljawad, S. Foulk, J., S. R. Kalidindi, and H. Lim. Application of the materials knowledge systems framework to establishing the linkage between microstructure and damage performance for two-phase composites. *Acta Mater. (under review)*, 2017.
- [70] J. Rice and D. Tracey. On the ductile enlargement of voids in triaxial stress fields. *Journal of the Mechanics and Physics of Solids*, 17(3):201 – 217, 1969.
- [71] J. R. Rice. Inelastic constitutive relations for solids, an internal-variable theory and its application to metal plasticity. *J. Mech. Phys. Solids*, 19:443–455, 1971.



- [72] T. M. Rodgers, J. D. Madison, V. Tikare, and M. C. Maguire. Predicting mesoscale microstructural evolution in electron beam welding. *JOM*, 68(5):1419–1426, 2016.
- [73] J. Ryan, L. Halpern, and M. Borrel. Domain decomposition vs. overset Chimera grid approaches for coupling CFD and CAA. *7th International Conference on Computational Fluid Dynamics, Big Island, Hawaii, ICCFD7*–1205, 7 2012.
- [74] A. Salinger. Albany: using agile components to develop a flexible, generic multiphysics analysis code. *International Journal for Multiscale Computational Engineering*, 14:415–438, 2016.
- [75] C. San Marchi, B. Somerday, D. Balch, and H. Jackson. Hydrogen-assisted deformation and fracture of austenitic stainless steels. 2nd International Conference Engineering Against Fracture (ICEAF II), Mykonos, Greece, 2011.
- [76] H. Schwarz. Über einen Grenzübergang durch alternierendes verfahren. *Vierteljahrsschrift der Naturforschenden Gesellschaft in Zurich*, 15:272–286, 1870.
- [77] J. Simo and T. Hughes. *Computational Inelasticity*. Springer-Verlag, New York, 1997.
- [78] J. Simo and C. Miehe. Associative coupled thermoplasticity at finite strains: Formulation, numerical analysis and implementation. *Computer Methods in Applied Mechanics and Engineering*, 98:41–104, 1992.
- [79] M. Staten. Cubit geometry and mesh generation toolkit 15.3 user documentation. Technical Report SAND2017-6895, Sandia National Laboratory, 2017.
- [80] W. Sun, J. Ostien, and A. Salinger. A stabilized assumed deformation gradient finite element formulation for strongly coupled poromechanical simulations at finite strain. *International Journal for Numerical and Analytical Methods in Geomechanics*, 37:2755–2788, 2013.
- [81] A. Taha and P. Sofronis. A micromechanics approach to the study of hydrogen transport and embrittlement. *Engineering Fracture Mechanics*, 68:803–837, 2001.
- [82] V. Tvergaard and A. Needleman. Analysis of the cup-cone fracture in a round tensile bar. *Acta metallurgica*, 32(1):157–169, 1984.
- [83] T. Werder, J. Walther, and P. Koumoutsakos. Hybrid atomistic-continuum method for the simulation of dense fluid flows. *Journal of Computational Physics*, 205:373–390, 2005.
- [84] Wolfram Research, Inc. Mathematica 11.0, 2017. URL <https://www.wolfram.com>.
- [85] Q. Yang, L. Stainier, and M. Ortiz. A variational formulation of the coupled thermo-mechanical boundary-value problem for general dissipative solids. *Journal of the Mechanics and Physics of Solids*, 54:401–424, 2006.
- [86] D. Zöllner. A new point of view to determine the simulation temperature for the potts model simulation of grain growth. *Computational Materials Science*, 86:99 – 107, 2014.

## A Need for concurrent, two-way coupling

Often we are asked to justify two-way coupling. If one is passionate about localization, we might wonder why one must justify that the near-field unloads the far-field. This is, in fact, a definition of localization. Experiments do just that. Whether the process be structural softening or material softening, two-way coupling is needed to ensure that the far-field and the near-field are in “balance.” Pontification aside, we wonder if one can construct a thought experiment that illustrates the need for two-way coupling. We envision the mechanical analog noted in Figure A.1 where the force in the coarse spring for one-way coupling  $P_c^1$  is

$$P_c^1 = P_{max} \exp(1) \left( \frac{\Delta_0}{L} \right) \quad (63)$$

while the force for two-way coupling  $P_c^2$  is

$$P_c^2 = P_{max} \exp(1) \left( \frac{\Delta_c}{L_c} \right). \quad (64)$$

The force in the fine scale spring is always

$$P_f = P_{max} \exp(1) \left( \frac{\Delta_f}{L_f} \right) \exp \left( - \frac{\Delta_f}{L_f} \right) \quad (65)$$

where the peak load at the fine scale is  $P_{max}$  is attained when  $\Delta_f = L_f$ . To be clear, we have intentionally added softening (read localization) into the fine scale through the exponential decay of the fine scale constitutive relation. If we only consider one-way coupling, the analytical solution for the displacement field in the coarse spring is

$$\Delta^1(x) = \left( \frac{\Delta_0}{L} \right) x \quad (66)$$


and the corresponding displacement field interpolated to the fine scale spring is

$$\Delta_f^1 = \left( \frac{\Delta_0}{L} \right) L_f. \quad (67)$$


Although idealized, this what the community terms sub-modeling. To emphasize the fact that one-way coupling is just a transfer of the coarse scale to the fine scale, we elect to employ the term “one-way transfer.” We solve the the far-field and interpolate the solution of the primary variable onto the near-field. Given the solution of the fine scale, we can calculate the force in the fine scale spring that results from one-way transfer,  $P_f^1$ , to be

$$P_f^2 = P_{max} \exp(1) \left( \frac{\Delta_0}{L} \right) \exp \left( - \frac{\Delta_0}{L} \right) \quad (68)$$

and note that the coarse-scale stiffness and fine-scale stiffness are equivalent at  $\Delta_0 = 0$ . Neat. To be clear, the end displacement of the fine scale  $\Delta_f$  does not depend on the fine-scale behavior. That



**Figure A.1.** Thought experiment to investigate the ramifications of one-way and two-way coupling. The figure is to scale with  $L = 5L_f$ .




**Figure A.2.** In one-way transfer, the coarse scale and the fine scale are not in equilibrium,  $P_c \neq P_f$ .

displacement is only governed by the coarse-scale solution. Consequently, there is no guarantee that the bodies will be in equilibrium,  $P_f^1 = P_c^1$ . In fact, examining Equation 63 and Equation 68, one finds that the bodies are never in equilibrium for finite  $\Delta_0$ . This property of one-way transfer is illustrated in Figure A.2. In contrast, two-way coupling ensures that the both the coarse scale and the fine scale are in equilibrium. Both scales are intimately connected through a degree of freedom,  $\Delta_f$ . To solve for equilibrium, we define a residual  $R$  as the difference between  $P_c^2$  and  $P_f^2$  knowing that  $L_c = 4L_f$

$$R = P_{max} \exp(1) \left[ \frac{\Delta_0 - \Delta_f}{4L_f} - \frac{\Delta_f}{L_f} \exp\left(-\frac{\Delta_f}{L_f}\right) \right]. \quad (69)$$

Employing Newton's method to solve the nonlinear equation  $R = 0$  for each  $\Delta_0$ , we contrast the two-way coupling solution to the one-way transfer in Figure A.3. For two-way coupling, we



**Figure A.3.** Two-way coupling localizes the displacement field in the fine scale. We remark that not only is one-way transfer inaccurate but the methodology is not conservative.

are plotting the load-displacement curve for the entire structure. For the one-way transfer, we are plotting the load-displacement curve for the fine scale. As expected, two-way coupling localizes and exhibits failure at lower values of  $\Delta_0$ . The coarse scale unloads to drive the localization of the displacement field at the fine scale. In contrast, one-way transfer does not localize the solution. The coarse scale load-displacement curve in Figure A.3 mirrors Figure A.2.

The fact that the fine-scale still retains substantial load-bearing capacity for one-way transfer illustrates that the methodology is not conservative. Clearly this example is idealized. For some boundary-value problems, differences between a one-way transfer and two-way coupling may be minimal. Not all boundary-value problems exhibit localization. Those differences, however, can only be quantified if one can employ two-way coupling.

## DISTRIBUTION:

1	MS 0897	Ted Blacker, 1543
1	MS 9042	Arthur Brown, 8259
1	MS 0840	Eliot Fang, 1554
1	MS 1411	Richard Muller, 1864
1	MS 9042	Scott Peterson, 8343
1	MS 1321	Veena Tikare, 1444
1	MS 0840	Jim Redmond, 1550
1	MS 9161	Christian Mailhiot, 8340
1	MS 0899	Technical Library, 8944 (electronic copy)





